

LAWRENCE
LIVERMORE
NATIONAL
LABORATORY

LLNL-TR-454593

Advanced Simulation and Computing Fiscal Year 2011-2012 Implementation Plan, Revision 0.5

M. McCoy, J. Phillips, C. Wampler, R. Meisner

September 13, 2010

Disclaimer

This document was prepared as an account of work sponsored by an agency of the United States government. Neither the United States government nor Lawrence Livermore National Security, LLC, nor any of their employees makes any warranty, expressed or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States government or Lawrence Livermore National Security, LLC. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States government or Lawrence Livermore National Security, LLC, and shall not be used for advertising or product endorsement purposes.

This work performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344.



FY11-FY12 Implementation Plan

Volume 2, Rev. 0.5

October 5, 2010

A large, stylized logo for ASC (Advanced Simulation & Computing) is centered on the page. It consists of the letters 'A', 'S', and 'C' in a bold, black, sans-serif font. The 'A' is formed by two thick, curved strokes that meet at the top. The 'S' is a single, thick, curved stroke. The 'C' is a thick, curved stroke that is open on the right side. To the right of the 'C' is a small 'TM' trademark symbol.

ASC™

LLNL-TR-454593

NA-ASC-120R-10-Vol.2-Rev.0.5-IP

This work performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344.

This document was prepared as an account of work sponsored by an agency of the United States government. Neither the United States government nor Lawrence Livermore National Security, LLC, nor any of their employees makes any warranty, expressed or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States government or Lawrence Livermore National Security, LLC. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States government or Lawrence Livermore National Security, LLC, and shall not be used for advertising or product endorsement purposes.



Advanced Simulation and Computing

FY11–12 IMPLEMENTATION PLAN

Volume 2, Rev. 0.5

October 5, 2010

Approved by:

Robert Meisner,

NNSA ASC Program Director

Signature

10/5/10

Date

Julia Phillips,

SNL ASC Executive

Concurrence Received

9/27/10

Michel McCoy,

LLNL ASC Executive

Concurrence Received

9/20/10

Cheryl Wampler,

LANL ASC Executive

(acting)

Concurrence Received

9/28/10

ASC Focal Point

Robert Meisner

NA-114

Tele.: 202-586-0908

FAX: 202-586-0405

bob.meisner@nnsa.doe.gov

IP Focal Point

Atinuke Arowojolu

NA-114

Tele.: 202-586-0787

FAX: 202-586-7754

atinuke.arowojolu@hq.doe.gov

Implementation Plan Contents at a Glance

Section No./Title		Vol. 1	Vol. 2
I.	Executive Summary		
II.	Introduction		
III.	Accomplishments		
IV.	Product Descriptions		
V.	ASC Level 1 and 2 Milestones		
VI.	ASC Roadmap Drivers for FY11–FY12		
VII.	Performance Measures and Data		
VIII.	Budget		
Appendix A.	Glossary		
Appendix B.	Codes		
Appendix C.	Points of Contact		
Appendix D.	Academic Alliance Centers		
Appendix E.	ASC Obligation/Cost Plan		

Contents

I. EXECUTIVE SUMMARY.....	1
II. INTRODUCTION.....	2
ASC Contributions to the Stockpile Stewardship Program	3
III. ACCOMPLISHMENTS FOR FY09–FY10.....	5
Computational Systems and Software Environment.....	5
Facility Operations and User Support.....	6
Academic Alliances	7
IV. PRODUCT DESCRIPTIONS BY THE NATIONAL WORK BREAKDOWN STRUCTURE.....	12
<i>WBS 1.5.4: Computational Systems and Software Environment</i>	<i>12</i>
WBS 1.5.4.1: Capability Systems	12
WBS 1.5.4.1 Systems Requirements and Planning (LANL).....	12
WBS 1.5.4.1 Red Storm Capability Computing Platform (SNL).....	13
WBS 1.5.4.1 Alliance for Computing at Extreme Scale Cielo Capability Computing Platform (SNL, LANL).....	14
WBS 1.5.4.1 Alliance for Computing at Extreme Scale Architecture Office (LANL, SNL)	14
WBS 1.5.4.2: Capacity Systems	15
WBS 1.5.4.2 Capacity Computing Planning and Integration (LLNL).....	16
WBS 1.5.4.2 Capacity System Integration (LANL)	16
WBS 1.5.4.2 ASC Capacity Systems (SNL).....	17
WBS 1.5.4.3: Advanced Systems.....	18
WBS 1.5.4.3 BlueGene/P and BlueGene/Q Research and Development (LLNL)	18
WBS 1.5.4.3 Petascale Application Enablement (LLNL).....	19
WBS 1.5.4.3 Sequoia (LLNL)	20
WBS 1.5.4.3 Hyperion Test Bed (LLNL).....	21
WBS 1.5.4.3 Advanced Systems Technology Research and Development (SNL).....	22
WBS 1.5.4.4: System Software and Tools	23
WBS 1.5.4.4 System Software Environment for Scalable Systems (LLNL)	24
WBS 1.5.4.4 Applications Development Environment and Performance Team (LLNL)	25
WBS 1.5.4.4 Software Support (LANL)	26
WBS 1.5.4.4 Applications Readiness (LANL).....	27
WBS 1.5.4.4 Productivity Project (LANL)	28
WBS 1.5.4.4 Programming Models for the Next-Generation Scientific-Computing Environment (LANL).....	29
WBS 1.5.4.4 Systems Research (LANL)	30
WBS 1.5.4.4 Data Intensive Supercomputing Project (LANL)	31
WBS 1.5.4.4 Application Enablement for Extreme-Scale Computing (LANL).....	32
WBS 1.5.4.4 Software and Tools for Scalability and Reliability Performance (SNL).....	33
WBS 1.5.4.4 System Simulation and Computer Science (SNL).....	35
WBS 1.5.4.5: Input/Output, Storage Systems, and Networking.....	36
WBS 1.5.4.5 Archive Storage (LLNL).....	36

WBS 1.5.4.5 Parallel and Network File Systems (LLNL)	38
WBS 1.5.4.5 Networking and Test Beds (LLNL)	39
WBS 1.5.4.5 File Systems and Input/Output Project (LANL)	40
WBS 1.5.4.5 Archival Storage Design and Development (LANL)	40
WBS 1.5.4.5 Archival Storage (SNL)	41
WBS 1.5.4.5 Scalable Input/Output Research (SNL)	42
WBS 1.5.4.6: Post-Processing Environments	43
WBS 1.5.4.6 Scientific Visualization (LLNL)	44
WBS 1.5.4.6 Scientific Data Management (LLNL)	45
WBS 1.5.4.6 Visualization and Insight for Petascale Simulations Project (LANL)	46
WBS 1.5.4.6 Production Systems for Visualization and Insight Project (LANL)	47
WBS 1.5.4.6 Physics-Based Simulation Analysis Project (LANL)	48
WBS 1.5.4.6 Scalable Data Analysis (SNL)	49
WBS 1.5.4.7: Common Computing Environment	50
WBS 1.5.4.7 Tripod Operating System Software (LLNL, LANL, SNL)	51
WBS 1.5.4.7 Open SpeedShop (LLNL, LANL, SNL)	52
WBS 1.5.4.7 Workload Characterization (LLNL, LANL, SNL)	53
WBS 1.5.4.7 Shared Work Space (LLNL, LANL, SNL)	54
WBS 1.5.4.7 Gazebo Test and Analysis Suite (LANL, SNL)	55
WBS 1.5.4.7 Debugger (LLNL, LANL, SNL)	56
WBS 1.5.4.7 Inter-Site High Performance Computing Security Integration (LLNL, LANL, SNL)	57
WBS 1.5.4.7 Open Source Contract Maintenance (LLNL, LANL, SNL)	58
WBS 1.5.4.7 OpenMPI Integration/Scaling (LLNL, LANL, SNL)	59
WBS 1.5.4.7 File System Architecture and Integration (LLNL, LANL, SNL)	60
WBS 1.5.4.7 SARAPE Re-Development (LLNL, LANL, SNL)	61
WBS 1.5.5: Facility Operations and User Support	63
WBS 1.5.5.1: Facilities, Operations, and Communications (Retired)	63
WBS 1.5.5.2: User Support Services	63
WBS 1.5.5.2 Hotlines and System Support (LLNL)	64
WBS 1.5.5.2 Integrated Computing Network Consulting, Training, Documentation, and External Computing Support (LANL)	64
WBS 1.5.5.2 Tri-Lab System Integration and Support (SNL)	65
WBS 1.5.5.2 User Support (SNL)	66
WBS 1.5.5.3: Collaborations	67
WBS 1.5.5.3 Program Support (LLNL)	68
WBS 1.5.5.3 Scientific Collaborations (LLNL)	68
WBS 1.5.5.3 Program Support (LANL)	69
WBS 1.5.5.3 Collaborations (SNL)	69
WBS 1.5.5.4: System and Environment Administration and Operations	70
WBS 1.5.5.4 System and Environment Administration and Operations (LLNL)	71
WBS 1.5.5.4 System Administration and Storage (LANL)	73
WBS 1.5.5.4 Operations and Procurement Support (LANL)	74
WBS 1.5.5.4 Requirements Planning (LANL)	75
WBS 1.5.5.4 Cielo Capability Computing Platform Initial Deployment (LANL)	76
WBS 1.5.5.4 Production Computing Services (SNL)	76
WBS 1.5.5.5: Facilities, Network, and Power	77
WBS 1.5.5.5 Facilities, Network and Power (LLNL)	78
WBS 1.5.5.5 Facilities and Power (LANL)	80
WBS 1.5.5.5 Networking (LANL)	81
WBS 1.5.5.5 Facilities, Networking, and Power (SNL)	82
V. ASC LEVEL 1 AND 2 MILESTONES	83

VI. ASC ROADMAP DRIVERS FOR FY11–FY12.....	106
VII. PERFORMANCE MEASURES.....	107
VIII. BUDGET.....	109
APPENDIX A. GLOSSARY.....	112
APPENDIX C. POINTS OF CONTACT	116
APPENDIX D. WBS 1.5.1.4-TRI-001 ACADEMIC ALLIANCE CENTERS.....	118
California Institute of Technology	118
Purdue.....	119
Stanford University	120
University of Michigan.....	122
University of Texas.....	124
APPENDIX E. ASC OBLIGATION/COST PLAN.....	125

List of Tables

Table V-1. ASC Level 1 <i>Proposed</i> Milestones and Interfaces with Defense Programs Components from FY12–FY16	83
Table V-2. Quick Look: Level 2 Milestone Dependencies for FY11	84
Table V-3. Quick Look: <i>Preliminary</i> Level 2 Milestone Dependencies for FY12.....	86
Table VII-1. ASC Campaign Annual Performance Results (R) and Targets (T).....	107

List of Figures

Figure E-1. ASC obligation/cost plan for FY11.	125
---	-----

I. Executive Summary

The Stockpile Stewardship Program (SSP) is a single, highly integrated technical program for maintaining the surety and reliability of the U.S. nuclear stockpile. The SSP uses past nuclear test data along with current and future non-nuclear test data, computational modeling and simulation, and experimental facilities to advance understanding of nuclear weapons. It includes stockpile surveillance, experimental research, development and engineering (D&E) programs, and an appropriately scaled production capability to support stockpile requirements. This integrated national program requires the continued use of current facilities and programs along with new experimental facilities and computational enhancements to support these programs.

The Advanced Simulation and Computing Program (ASC) is a cornerstone of the SSP, providing simulation capabilities and computational resources to support the annual stockpile assessment and certification, to study advanced nuclear weapons design and manufacturing processes, to analyze accident scenarios and weapons aging, and to provide the tools to enable stockpile Life Extension Programs (LEPs) and the resolution of Significant Finding Investigations (SFIs). This requires a balanced resource, including technical staff, hardware, simulation software, and computer science solutions.

In its first decade, the ASC strategy focused on demonstrating simulation capabilities of unprecedented scale in three spatial dimensions. In its second decade, ASC is focused on increasing its predictive capabilities in a three-dimensional (3D) simulation environment while maintaining support to the SSP. The program continues to improve its unique tools for solving progressively more difficult stockpile problems (focused on sufficient resolution, dimensionality, and scientific details); to quantify critical margins and uncertainties; and to resolve increasingly difficult analyses needed for the SSP. Moreover, ASC has restructured its business model from one that was very successful in delivering an initial capability to one that is integrated and focused on requirements-driven products that address long-standing technical questions related to enhanced predictive capability in the simulation tools.

ASC must continue to meet three objectives:

- **Objective 1. Robust Tools.** Develop robust models, codes, and computational techniques to support stockpile needs such as refurbishments, SFIs, LEPs, annual assessments, and evolving future requirements.
- **Objective 2. Prediction through Simulation.** Deliver validated physics and engineering tools to enable simulations of nuclear weapons performance in a variety of operational environments and physical regimes and to enable risk-informed decisions about the performance, safety, and reliability of the stockpile.
- **Objective 3. Balanced Operational Infrastructure.** Implement a balanced computing platform acquisition strategy and operational infrastructure to meet Directed Stockpile Work and SSP needs for capacity and high-end simulation capabilities.

II. Introduction

The ASC Program supports the National Nuclear Security Administration's (NNSA's) overarching goal of nuclear weapons stewardship: *"We continue to advance the Stockpile Stewardship Program to push the scientific and engineering boundaries needed to maintain our nuclear arsenal. It also means maintaining the basic science and engineering that is the foundation of the weapons program."*¹

In 1996, ASCI—the Accelerated Strategic Computing Initiative—was established as an essential element of the SSP to provide nuclear weapons simulation and modeling capabilities.

In 2000, the NNSA was established to carry out the national security responsibilities of the Department of Energy (DOE), including maintenance of a safe, secure, and reliable stockpile of nuclear weapons and associated materials capabilities and technologies.

Shortly thereafter, in 2002, ASCI matured from an initiative to a recognized program and was renamed the ASC Program.

Prior to the start of the nuclear testing moratorium in October 1992, the nuclear weapons stockpile was maintained through 1) underground nuclear testing and surveillance activities and 2) "modernization" (in other words, development of new weapons systems). A consequence of the nuclear test ban is that the safety, performance, and reliability of U.S. nuclear weapons must be ensured by other means for systems far beyond the lifetimes originally envisioned when the weapons were designed.

NNSA will carry out its responsibilities through the twenty-first century in accordance with the current Administration's vision and the Nuclear Posture Review (NPR) guidance. NNSA Administrator Thomas P. D'Agostino summarized² the NNSA objectives for the SSP as follows:

"Our fundamental national security responsibilities for the United States include:

- *Assuring the safety, security and reliability of the U.S. nuclear weapons stockpile while at the same time transforming the stockpile and the infrastructure that supports it;*
- *Reducing the threat posed by nuclear proliferation; and,*
- *Providing reliable and safe nuclear reactor propulsion systems for the U.S. Navy."*

"Throughout the past decade, the Stockpile Stewardship Program has proven its ability to successfully sustain the safety, security and reliability of the nuclear arsenal without resorting to underground nuclear testing. The SSP also enables the U.S. to provide a credible strategic deterrent capability with a stockpile that is significantly smaller."

Additionally, key investment recommendations cited in the 2010 NPR are to:

- Strengthen the science, technology, and engineering base needed for conducting weapon system LEPs
- Mature advanced technologies to increase weapons surety

¹ NNSA Strategic Planning Guidance for FY2010–2014, April 2008, page 17.

² Testimony on the FY 2008 National Defense Authorization Budget Request for the Department of Energy's NNSA before the House Armed Services Subcommittee, March 20, 2007.

- Qualify weapon components and certify weapons without nuclear testing
- Provide annual stockpile assessments through weapons surveillance

This strategy includes developing and sustaining high-quality scientific staff, as well as supporting computational and experimental capabilities.³

The ASC Program plays a vital role in the NNSA infrastructure and its ability to respond to the NPR guidance. The program focuses on developing modern simulation tools that can provide insights into stockpile problems, providing tools with which designers and analysts can certify nuclear weapons, and guiding any necessary modifications in nuclear warheads and the underpinning manufacturing processes. Additionally, ASC is enhancing the predictive capability necessary to evaluate weapons effects, designing experiments, and ensuring test readiness.

ASC continues to improve its unique tools to solve progressively more difficult stockpile problems—with a focus on sufficient resolution, dimensionality, and scientific details—to enable quantification of margins and uncertainties (QMU) and to resolve the increasingly difficult analyses needed for stockpile stewardship. The DSW program provides requirements for simulation. These requirements include planned LEPs, stockpile support activities that may be ongoing or require short-term urgent response, and requirements for future capabilities to meet longer-term stockpile needs. ASC's advanced, leading-edge technologies in high performance computing (HPC) and predictive capability meet DSW's short- and long-term needs, including the annual assessments and certifications and SFIs.

The following section provides an overview of FY10 and upcoming ASC contributions to the SSP:⁴

ASC Contributions to the Stockpile Stewardship Program

In FY10, ASC continued to deliver science-based simulation tools to support annual assessments and the next generation of LEPs. Code suite physics and optimization were completed in support of the National Technical Nuclear Forensics program and High Energy Density Physical experimental program. The Predictive Capability Assessment Project began under the Verification and Validation (V&V) program element to measure and monitor progress toward a physics-based simulation capability that does not rely on calibration. An initial assessment of new capabilities in a primary burn code was performed. ASC also provided tools for both experiment and diagnostic design to support the indirect-drive ignition experiments on the National Ignition Facility (NIF) and for improved confidence and response time for questions of vital importance to achieving predictive simulation capability. In addition, ASC continued to provide national leadership in HPC and deploy capability and capacity platforms in support of Defense Programs campaigns. Roadrunner, the advanced architecture petaFLOPS hybrid high-performance computer was formally transitioned to production computing for weapons science applications.

In FY 11, ASC will continue delivering science-based simulation tools for annual assessments and next-generation LEPs, focusing on improved physics, fidelity, and calculations in support of DSW and the National Code Strategy. The methodology for predictive capability assessment will be demonstrated in FY11 for a limited set of

³ 2010 Nuclear Posture Review Report, April 2010, p. 42.

⁴ FY11 ASC Program Plan for past, present, and planned contributions to the SSP.

simulations common to both physics laboratories. The ability to simulate full system performance near thresholds where data are sparse will be assessed. Cielo, the next generation ASC National User Facility, will be in operation, replacing Purple to provide capability computing cycles for the SSP. Development of the advanced architecture Sequoia high-performance computer will continue, with a focus on Scalable Applications Preparation (SAP) and outreach. Installation and operation of the next-generation tri-lab Linux capacity clusters (TLCC) and the associated common user environment will continue across the three NNSA laboratories.

In FY12 and beyond, ASC will be focusing on strengthening the science basis and driving down uncertainties for weapons simulations to a degree that NNSA can ultimately, and credibly, claim predictive capability; instituting a robust, formalized peer review system; increasing the number of production computing cycles to support increased use of simulation in stockpile activities and reliance on uncertainty quantification (UQ) in weapons decisions; and pursuing exascale computing to meet time-urgent, future capability needs as documented in the *ASC Roadmap* and the *Predictive Capability Framework* (PCF).

III. Accomplishments for FY09–FY10

ASC accomplishments from Quarter 4, fiscal year 2009, through quarter 3, fiscal year 2010, are reflected below for the Computational Systems and Software Environment (CSSE) and Facility Operations and User Support (FOUS) sub-programs.

ASC headquarters (HQ) is pleased to highlight the outstanding achievements of the Defense Programs laboratories.

Computational Systems and Software Environment

LLNL Accomplishments for Computational Systems and Software Environment

- Performed advanced application preparation work for Sequoia and demonstrated capabilities through science runs before transition of the Sequoia Initial Delivery (ID) System (Dawn) to the secure computing facility (SCF)
- Deployed and supported the Sequoia hardware environment, including file systems, archival storage, visualization clusters, and networking infrastructure, as well as the software environment
- Delivered Capability Computing Campaign (CCC) capabilities to the ASC Program via Purple
- Planned for the next capacity tri-lab procurement and common computing software environment and supported the systems software and tools on the current systems

LANL Accomplishments for Computational Systems and Software Environment

- Successfully completed Roadrunner Phase 3 transition to operational status
- Successfully completed Critical Decision (CD) 4b for Roadrunner Phase 3
- Awarded Cielo contract to Cray Inc.
- Developed testing strategies for Cielo file system and interconnect

SNL Accomplishments for Computational Systems and Software Environment

- Received acceptance from the Linux community for the research project to investigate general-purpose computing on graphics processing units (GPGPU) redundant array of independent disks (RAID) and began plans to transition to a production product; the success of the RAID project also led to a collaboration with LSI to design and develop next-generation RAID controllers
- Released the initial implementation of tools for analysis of ensembles of runs, promoting interactive investigation of entire sets of runs as a single dataset

Tri-Lab Accomplishments for Computational Systems and Software Environment

- Deployed Common Computing Environment (CCE) capabilities across all TLCC systems, including Open | SpeedShop (O | SS), Shared Work Space, Gazebo Test and Analysis Suite, configuration management of the Tripod Operating System Software (TOSS), and performance monitoring tools

Facility Operations and User Support

LLNL Accomplishments for Facility Operations and User Support

- Integrated the Sequoia ID (Dawn) system into the classified environment and brought the system into General Availability (GA) status. Also, the Interactive Data Analysis (IDA) cluster was integrated into the classified environment and the system brought into production status. LLNL tracked and placed contracts and licenses needed for system operations and vendor support; performed ongoing hardware self-maintenance of BlueGene/L, parallel global file systems, Peloton, and TLCC capacity systems. LLNL expanded its deployment of an identity management solution to support electronic workflows for approvals, management, and provisioning of groups. In addition, LLNL deployed in the unclassified environment a one-time password authentication mechanism using standards-based protocols and upgraded the security infrastructure to operate under the NNSA Policy Letters (NAPS)-compliant *Livermore Computing (LC) Information System Security Plan*. The first phase of LLNL's trouble ticket system replacement was completed. The FrontRange incident management module replaced the Remedy trouble ticket system.
- Completed the east room 7.5-MW electrical distribution expansion from the new electrical yard into the first level of the machine room and fully commissioned and completed activation testing of the electrical components from the first phase of construction of the B-453 15-MW power expansion (7.5MW west room electrical equipment).
- Completed and implemented energy savings initiatives resulting from self-benchmarking tools and computational fluid dynamic analysis for B-453 (the Terascale Simulation Facility, TSF). To date the temperatures in the computer rooms and on the chilled water supply have been raised for an annual energy savings of over 47,500,000 kwh/yr. These energy saving initiatives assisted in B-453 achieving U.S. Green Building Council LEED Gold certification on December 24, 2009.

LANL Accomplishments for Facility Operations and User Support

- Successfully completed the Infrastructure Equipment Upgrade Project in preparation for the Cielo supercomputer. The upgrade increased power at the Nicholas C. Metropolis Center for Modeling and Simulation (SCC) to 19.2 MW and cooling to 6,000 tons. It also included the installation of a heat exchanger system in preparation for future liquid-cooled computers.
- Prepared the site for the installation of Cielo, Panasas for Cielo, and the Cielo testbed.
- Deployed monitoring infrastructure to all TLCC- and Roadrunner-class platforms; monitoring interface and function adapted to first-level hardware maintenance support.
- Decommissioned the Lightning/Bolt and Flash/Gordon LNXI systems in preparation for future HPC platforms at the SCC.
- Successfully passed the Return Material Authorization security audit, which included a review of procedures for returning failed parts to the appropriate computer vendors. These procedures ensure there is a complete audit trail for the de-energizing of failed components needed prior to leaving the SCC.
- Completed initial phase of electrical metering project to remotely monitor electrical power used by the high-performance computers and more accurately measure

Power Usage Effectiveness (PUE) at the SCC in an effort to optimize energy efficiencies and identify cost savings.

- Added six PaScaLBB lanes to existing input/output (I/O) infrastructure, which brings total file system I/O bandwidth to 900 GB/sec. Configuration of the twelve lanes provided connectivity to all HPC clusters in the secure network.
- Developed user training course for Cielo.

SNL Accomplishments for Facility Operations and User Support

- Transitioned Red Storm to support Other Government Agency (OGA) opportunities per NNSA/ASC guidance
- Achieved GA of TLCC capacity systems Whitney (SNL/CA), Unity (SNL/NM), and Glory (SNL/NM)
- Increased Wide Area Network (WAN) reliability by completing fully redundant fiber paths, eliminating one single run connection that remained in California

Academic Alliances

University of Chicago Accomplishments

- FLASH, the highly capable, fully modular, extensible AMR community code developed by the Flash Center, has now been used by more than 700 scientists around the world, and the results of more than 375 papers directly use FLASH.
- Won a DOE award for the Flash Center Early Science Proposal for Mira, the 10 petaFLOPS IBM BG/Q computer that will be deployed at the Argonne Leadership Computing Facility in 2012. Included in the award is assistance by personnel at Argonne National Laboratory to help the FLASH code perform at scale on Mira plus substantial computing resources, including 5 M CPU-hours per year for this purpose and 150 M CPU-hours for early science.
- Initiated a major effort to add capabilities to the FLASH code to make it an open toolset for the academic high-energy density community. The capabilities added this year include un-split PPM hydrodynamics solver, Spitzer conductivity for hydrodynamics and MHD, super time stepping algorithm that mitigates the disparate time scales between diffusion and advection processes, three temperature treatment of electrons, ions, and radiation, ray tracing and laser energy deposition, and Chombo patch-based mesh.
- Produced significant progress in Flash Center simulations of Type Ia supernovae (SNe Ia) for understanding these explosions. Within the past year, the Center showed the rate of burning in buoyancy-driven turbulent nuclear combustion appears to be governed by the length scale corresponding to the flame polishing length. This suggests a treatment of such combustion that can capture the burning that occurs at length scales smaller than those that can be resolved in whole-star simulations of SNe Ia. The Center confronted a set of simulations of the gravitationally confined detonation model of SNe Ia with high-quality observations and found that the light curves predicted by the simulations match observed light curves reasonably well. Finally, the Center demonstrated that viewing the asymmetric explosions produced by the model of SNe Ia from different directions may contribute significantly to the anomalous scatter in the calibration of these events as "standard candles" for determining the properties of dark energy.

University of Illinois at Urbana-Champaign Accomplishments

- Used *Rocpack* module to better understand the random packing of HE crystals
- Implemented new multiphysics framework to augment *Sierra*
- Extended and validated two CFD modules—*RocfluND* (unstructured, nondissipative code) and *RocfloCM* (employing Chimera meshing strategies)
- Showed how NASA-sponsored Constellation University Initiative Program research results explained erosive burning from first principles
- Finished NASA STTR project exploring the use of CSAR-technology in reduced-order clustering for UQ
- Finished AF-Edwards and AF-Eglin SBIR/STTR projects in tomographic analysis of energetic materials
- Finished Army-Redstone SBIR/STTR projects in cook-off and insensitive munitions impact
- Provided simulation software development and support, and engineering analysis services to government mission agencies and U.S. industry through the University of Illinois' spinoff company, IllinoisRocstar LLC, which has grown rapidly

University of Utah Accomplishments

- Extended Uintah to use new dynamic task scheduling and load balancing algorithms. Together with a new tile-based adaptive mesh refinement (AMR) algorithm, this has made it possible for Uintah to scale and to routinely run on 98K cores on NSF machines such as Kraken at NICS.
- Performed two demonstrations of quantifying uncertainty of velocities and species concentrations in non-reacting turbulent plumes (He plumes) and quantifying uncertainty of heat flux to a steel container in open JP8 pool fires. Datasets were provided by the FLAME facility from SNL. These two problems represent scalable V&V/UQ for multi-physics, multi-scale problems.
- Extended the Arches component to include a reacting disperse multiphase representation using the direct quadrature method of moments and applied it to reacting coal gasification and oxy-firing. This represents one of the first large eddy simulation tools in the world with this capability.
- Extended the Uintah combustion computational model for PBX to simulate detonation and the deflagration-to-detonation transition. A paper describing the application of this technology to Hazard Analysis recently won the best paper award at the TeraGrid10 supercomputing conference.
- Extended the reaction chemistry model of JP8 fuel, the Utah Surrogate Mechanism, to include major components in gasoline beyond the reference fuels of normal heptane and isooctane. The mechanism was coupled with a University of Cambridge code to predict ignition delay, heat release, and soot formation in a gasoline engine. The resulting publication by a multinational (U.S., UK, and Japan) team was selected by Fédération Internationale des Sociétés d'Ingénieurs des Techniques de l'Automobile to be republished in German in its flagship journal MTZ Worldwide and ATZ auto technology.

California Institute of Technology (Caltech) Accomplishments

- Completed two full UQ yearly cycles in the ballistic range based on Caltech's "data-on-demand" UQ protocol. This protocol uses Concentration-of-Measure inequalities

to bound the probability that the system will not performed within specs. The systems considered in the runs consisted of: Steel projectiles and steel plates (Year 1); steel projectiles on steel plates (Year 2). A full set of uncertainties was determined in both cases, including modeled system uncertainties (uncertainties in the behavior of the system as predicted by the model), modeling-error uncertainties (uncertainties due to discrepancies between model and experiment), and experimental-scatter uncertainties. The steel/steel system exhibited very large uncertainties due to large experimental scatter and could not be certified with any degree of confidence. By contrast, the lethality of the steel/aluminum system could be certified with very high confidence factors.

- Developed, verified, deployed and tested a Lagrangian, message passing interface (MPI)-based, parallel Optimal Transportation Meshfree code and used it to compute to compute modeled system uncertainties in the Year-2 UQ campaign. The code has been deployed in the in-house CACR shc platform and in a number of NNSA laboratory platforms (including Coyote, uBG/L, Lobo, Hera, and Cerrillos) and exhibits good scalability properties. The code has also been interfaced with the UQ pipeline and proved highly predictive (as measured by Year-2 model-error diameters) in ballistic simulations.
- Developed, deployed, and tested a UQ pipeline that orchestrates the calculation of modeled system uncertainties (uncertainties in the behavior of the system as predicted by the model). Pyre, the Caltech ASC component framework, forms the architectural backbone of the UQ pipeline. The UQ pipeline provides capability for staging, monitoring and analyzing thousands of simulations in a heterogeneous and distributed computing mode that makes effective use of multiple NNSA platforms simultaneously. It also provides a suite of global optimizers for the computation of diameters.
- Developed and integrated into Caltech's Virtual Testing Facility, a finite-difference based solver for hyperelastic and viscoplastic systems using a variant of the weighted essentially non-oscillatory schemes popular in gas dynamics to solve the equations of motion expressed in an Eulerian formulation. The formulation allows for a wide range of constitutive relations. The formulation also includes a simple reflection technique combined with ghost-cells to enforce fixed boundaries with a zero tangential stress condition (such as, free slip). The approach has been verified based on 1D and 2D examples and using a range of constitutive laws with and without additional plastic modeling.
- Developed optimal probability-of-failure bounds that deliver the best possible bound given the information available about the system. The bounds computation of the bounds reduces to a finite-dimensional optimization problem. In some cases of practical interest, the optimal bounds can be determined explicitly. We have also developed a "legacy UQ protocol" that enables rigorous certification based on legacy—or archival—data. The protocol requires *a priori* knowledge of certain properties of the response function that need to be estimated by independent means.
- Completed the deployment into the SPHIR facility of a Coherent Verdi-V6 and Valyn VISAR velocity interferometer system for arbitrary reflectors. Completed the deployment and testing of capture media for 3D debris-cloud analysis. Completed the development of a Mach lens experimental technique that provides access to equation of state (EOS) data under conditions and for materials of interest to the hypervelocity impact application. Validation experiments of the Mach lens configuration have been carried out at Caltech and the STAR facility at SNL-NM.

Purdue University Accomplishments

- Performed fluid damping computations on the Center for Prediction of Reliability, Integrity, and Survivability of Microsystems (PRISM) device for Knudsen numbers in the slip regime. Uncertainty bars were computed on the damping simulations based on measured uncertainties in inputs obtained from metrology experiments at Purdue. Good comparison was found with damping measurements made at Purdue.
- Developed and implemented in MEMOSA coupled algorithms for the interaction of structural deformation, fluidic damping, and electrostatics force calculation; verification problems were computed.
- Developed a dielectric charging model that accounts for tunneling, emission, capture, and drift of carriers, and was implemented in MEMOSA-FVM. A number of verification tests were completed, and charging in a MIM capacitor was computed. Furthermore, first-principles atomistic simulations were initiated whose objective is to characterize defect energy levels that appear in the dielectric charging model.
- Performed atomistic contact simulations for metal-dielectric contact to determine surface hardness, which governs the dynamics of contact closing, and the pull-out stress, which governs contact opening. In addition, a mesoscale Hertzian contact model was developed for use in the device-level contact simulations.
- Developed a coarse-grained model of the complete PRISM device, including the structural response, electrostatic actuation, fluid damping, dielectric charging, creep, and a preliminary simplified model of membrane contact with the dielectric layer. Sensitivity computations were performed to identify the most important inputs and model parameters affecting pull-in and pull-out voltage and device lifetime.

Stanford University Accomplishments

- Developed new framework for full system QMU with multiple gates. Performed full system QMU with low-fidelity heat release model and flight conditions uncertainties estimated through Bayesian inversion.
- Implemented and verified new adjoint solver for compressible Reynolds-Averaged Navier Stokes (RANS) equations and performed preliminary tests of adjoint based goal oriented mesh adaptation and numerical error estimation.
- Expanded C++ computational infrastructure supporting system level integration using the “Joe” RANS and the high fidelity “Charles” large eddy simulation solvers to include local or global grid refinement to facilitate adjoint-based mesh adaptation for solution verification in Joe and grid sensitivity studies for Charles. Extended infrastructure’s post-processing capabilities to support scalable parallel dimensional reduction operations on the large datasets resulting from UQ ensemble simulations or transient records of individual simulations. Initiated several non-reacting and reacting simulation campaigns on sub-system and component level, including non-reacting and reacting jet in cross-flow, high-temperature gas dynamics laboratory (HTGL) reflected-shock/boundary layer interaction, unstart by mass injection experiment.
- Further developed and refined in-house experiments, including mixing and combustion in jet in supersonic cross-flow experiment, new high-fidelity H_2/O_2 reaction mechanism, unstart by mass injection, reflected shock/boundary layer interaction experiment, to serve as validation basis for different solvers and to better understand key physical aspects.

- Implemented minimal version of RANS solver Joe in new domain-specific language (DSL) Liszt with matrix support, allowing run on multiple back-ends. Organized first tutorial and dissemination event to collect feedback from different users.

University of Michigan Accomplishments

- Performed experiments to quantify the behavior of the Be plasma that drives the radiative shock.
- Modified the BATSRUS/CRASH code to: 1) include electron heat conduction and separate electron and ion temperatures; 2) function with radiation and AMR in 1D, 2D, and 3D; 3) include multigroup radiation diffusion; and 4) determine the multigroup opacities and other parameters needed for the previous items.
- Improved PDT's utility for high-fidelity rad-transfer modeling, by: 1) speeding up the single-core "grind time" by a factor of 30; 2) improving and testing scaling out to 12,288 cores on Hera; 3) implementing adaptive time-step controls; 4) implementing and applying electron energy sources; 5) creating and beginning to study CRASH-like rad-transfer test problems; 6) implementing new I/O model for large core counts; and 7) restructuring output to support improved visualization.
- Completed first predictive study and submitted papers for publication describing: 1) development of the physics-informed emulator needed in this study; 2) the approach to assessing predictive capability; and 3) the results of this study.
- Implemented numerous additional verification tests.

University of Texas, at Austin, Accomplishments

- Completed first forward propagation of uncertainty in the full system simulation to quantify the uncertainty in the ablation rate and peak heat flux
- Applied full validation cycle to model problems in combustion chemistry and turbulence
- Designed and implemented a library for method of manufacture solutions for many classes of problems
- Modeled the Electric Arc Shock Tube experiments of NASA to calibrate radiation and non-equilibrium chemistry models
- Developed and implemented parallel delayed rejection adaptive Markov chain Monte Carlo sampling algorithms for Bayesian inversion into our toolkit for UQ (QUESO)

IV. Product Descriptions by the National Work Breakdown Structure

The National Work Breakdown Structure divides ASC into five subprograms. Physics and Engineering Models, Integrated Codes, V&V are sub-programs presented in Volume 1 of the Implementation Plan. CSSE and FOUS are presented in Volume 2 of the Implementation Plan.

WBS 1.5.4: Computational Systems and Software Environment

The mission of this national sub-program is to build integrated, balanced, and scalable computational capabilities to meet the predictive simulation requirements of NNSA. This sub-program strives to provide users of ASC computing resources a stable and seamless computing environment for all ASC-deployed platforms, which include capability, capacity, and advanced systems. Along with these powerful systems that ASC will maintain and continue to field, the supporting software infrastructure that CSSE is responsible for deploying on these platforms includes many critical components, from system software and tools, to I/O, storage and networking, to post-processing visualization and data analysis tools, and to a CCE. Achieving this deployment objective requires sustained investment in applied R&D activities to create technologies that address ASC's unique mission-driven need for scalability, parallelism, performance, and reliability.

WBS 1.5.4.1: Capability Systems

This level 4 product provides capability production platforms and integrated planning for the overall system architecture commensurate with projected user workloads. The scope of this product includes strategic planning, research, development, procurement, hardware maintenance, testing, integration and deployment, and quality and reliability activities, as well as industrial and academic collaborations. Projects and technologies include strategic planning, performance modeling, benchmarking, and procurement and integration coordination. This product also provides market research for future systems.

Capability Systems Deliverables for FY11

- Production cycles of Purple until the machine is retired this year
- Level 2 integration milestone for Cielo, which includes acceptance testing of phases 1 and 2, security accreditation, and phase 1 and 2 systems in use for CCCs

WBS 1.5.4.1 Systems Requirements and Planning (LANL)

The Systems Requirements and Planning project covers all aspects of program and procurement planning for current and advanced systems and strategic planning for

supporting infrastructure. The major focus is to define requirements and potential system architectures for advanced systems platforms that meet ASC programmatic requirements and drivers. Additionally, this project provides a focus for the various planning efforts. In FY11, this project will focus on the project management of the Cielo system.

In FY10, LANL completed the ASC Level 2 milestone for operational readiness for Roadrunner, transitioned Roadrunner to operations, and completed CD4b for Roadrunner.

Planned activities in FY11:

- Exercise Cielo Phase-2 option
- Develop and implement CCC process for Cielo
- Provide oversight of Cielo system integration and deployment

Expected deliverables in FY11:

- Cielo ASC Level 2 System Integration milestone
- Acceptance testing of Cielo phases 1 and 2
- Security accreditation for Cielo

Preliminary planned activities in FY12:

- Complete Cielo ASC Level 2 Production Capability Readiness milestone
- Complete Cielo CD4
- Transition to operations

WBS 1.5.4.1 Red Storm Capability Computing Platform (SNL)

This project covers on-going activities for the Red Storm contract, funded from FY10 carry-over dollars. The maintenance-only contract ends in August 2011, and its outstanding activity involves migrating capability users.

Cray provides hardware support and software upgrades, including installation for the system. The contract calls for 24x7 response to outages. Support personnel are on-site from 7 a.m. to 5 p.m. During off hours, work is begun within two hours of contact and performed until complete.

In FY10, SNL delivered system modifications to support extended file auditing capabilities to meet the programmatic needs of some OGAs. The original Red Storm contract, which was placed with Cray near the end of FY02, was closed in August 2010 and replaced with a maintenance-only contract.

Planned activities in FY11:

- Continue production computing and complete transfer of resource to OGA projects (migration of capability users)

Expected deliverables in FY11:

- Transition of Red Storm capability work to Cielo

Preliminary planned activities in FY12:

- Use Red Storm for OGA projects

WBS 1.5.4.1 Alliance for Computing at Extreme Scale Cielo Capability Computing Platform (SNL, LANL)

The Cielo capability computing platform is being acquired and deployed under the New Mexico Alliance for Computing at Extreme Scale (ACES). ACES is a joint collaboration between LANL and SNL defined under a Memorandum of Understanding to provide a user facility for capability computing to the NNSA weapons programs in support of stockpile stewardship, to develop requirements and system architecture for ASC capability systems requirements definition, architecture design, procurement, key technology development, systems deployment, operations, and user support. A joint design team has developed tri-lab requirements and architectural specification for Cielo with input from the user communities at all three labs. The architecture and design of Cielo will be optimized to provide performance at the full scale of the machine, in support of the NNSA program's most challenging CCCs.

This project covers all aspects of program and procurement planning for current and future capability systems and strategic planning for supporting infrastructure. A major effort is to continue to identify tri-lab capability requirements.

In FY10, after a competitive procurement process, ACES awarded a contract to Cray, Inc. for the Cielo capability computing platform.

Planned activities in FY11:

- Prepare for Cielo Phase 1 and Phase 2 deliveries
- Complete Cielo Phase 1 and Phase 2 integration into the Strategic Computing Complex at LANL
- Complete Application Readiness activities in support of initial CCCs
- Transition Cielo platform into initial and full support of ASC CCCs

Expected deliverables in FY11:

- System Integration Readiness
- Acceptance of the Phase 1 Cielo System
- Application Readiness
- Acceptance of the integrated Phase 1 and Phase 2 systems
- Security accreditation
- Phase 1 CCC capability
- Phase 2 CCC capability

Preliminary planned activities in FY12:

- Prepare for Production Capability Readiness
- Complete CD4 for Cielo

WBS 1.5.4.1 Alliance for Computing at Extreme Scale Architecture Office (LANL, SNL)

The primary objective for the ACES architecture office is to define requirements and potential system architectures for platforms that meet future ASC programmatic requirements and drivers. Additionally, this project provides a focus for the various

planning efforts, such as the DOE Exascale Initiative, and provides project management support for those efforts.

The ACES architecture office will coalesce mission requirements, application algorithms, user requirements, and HPC computer industry hardware/software trends into the design process. When vendor designs/proposals are available, for example, in response to the Cielo request for proposal (RFP), equivalent analysis is required to identify and select a proposal that matches pre-established design criteria. The ACES architecture office will also identify, in collaboration with the computer industry, critical technology gaps for future production capability systems, and support new technology development projects in the new industry partnership program to be named FastForward, to address these issues. The FastForward projects are targeted to start in FY12 or possibly FY13.

In FY10, LANL and SNL began identifying technology trends that are applicable to the design of upcoming ASC capability systems and also provided technical oversight for the interconnection network D&E project that was awarded to Cray.

Planned activities in FY11:

- Support definition of the new baseline NNSA/ASC platform roadmap; clarify and define roles and interfaces between this ASC platform roadmap and the Exascale Initiative timeline
- Begin development of CD0 for ASC's next-generation capability system that will be the eventual replacement for the Cielo platform
- Extend the DOE Exascale Initiative partnership activities to monitor technology trends from the U.S. computer industry to understand and quantify the technology foundation for ASC's next-generation capability system and future ASC system architectures
- Lead the technical activities within ACES to support the collaboration with Cray on the interconnection network project
- Develop/define a systematic the codesign process for HPC

Expected deliverables in FY11:

- Support for the creation of the new NNSA/ASC platform roadmap (while we expect this deliverable to be due in FY11, responsibility for this will reside with NNSA/ASC and the tri-lab CSSE/FOUS program element leads)
- Published definition of MPI Event Trace format for Cray INP D&E

Preliminary planned activities in FY12:

- Obtain CD0 approval for ASC's next-generation capability system platform

WBS 1.5.4.2: Capacity Systems

This level 4 product provides capacity production platforms commensurate with projected user workloads. The scope of this product includes planning, research, development, procurement, hardware maintenance, testing, integration and deployment, and quality and reliability activities, as well as industrial and academic collaborations. Projects and technologies include the procurement and installation of capacity platforms.

Capacity Systems Deliverables for FY11

- Tri-lab TLCC2 Linux clusters to meet NNSA's capacity computing needs
- Installation and integration of capacity systems on the Scalable Unit (SU) model to meet programmatic requirements for resources into FY12 and beyond
- Integration of CCE selected stack and tools on TLCC clusters

WBS 1.5.4.2 Capacity Computing Planning and Integration (LLNL)

The LLNL ASC strategy for capacity computing is to leverage industry advances and open source software standards to build, field, and integrate Linux clusters of various sizes into classified and unclassified production service. The programmatic objective is to dramatically reduce overall total cost of ownership of these "capacity" systems relative to best practices in Linux cluster deployments today. This objective strives to quickly make these systems robust, useful production clusters under the coming load of ASC scientific simulation capacity workloads.

A contract will be placed for a large amount of capacity computing resources to be delivered at all three sites over the next two fiscal years. This tri-lab procurement will include options to procure a specific number of SUs each quarter from 3QFY11 through 3QFY12 for a total commitment of 32–50 SUs over four quarters. The procurement will include Level 3 support and an onsite parts cache for each site.

This contract will provide a substantial increase in capacity compute resources for NNSA at all three labs. Clusters will have a common architecture and run the CCE software stack at each site. Options for clusters with GPGPU technology have been accommodated per request from LANL.

In FY10, LLNL monitored the computing industry developments for opportunities to augment capacity computing and the associated infrastructure. LLNL also led the tri-lab in planning for the TLCC2 project.

Planned activities in FY11:

- Receive, install, and integrate the TLCC2 SUs assigned to LLNL
- Ensure complete system delivery and ensure maintenance is included

Expected deliverables in FY11:

- TLCC2 SUs delivered to LLNL operating in production computing environment running the TOSS stack
- TLCC1 systems at LLNL operating in production computing environment running the TOSS stack

Preliminary planned activities in FY12:

- Take delivery of additional SUs in FY12 for capacity computing
- Install and integrate the additional SUs assigned to LLNL

WBS 1.5.4.2 Capacity System Integration (LANL)

The Capacity System Integration project is responsible for integrating new ASC capacity systems into LANL classified, unclassified, and open computing networks. Included in this project is support for the ASC capacity system acquisition strategy along with

providing requirements that help to achieve the strategy. In FY11, the main focus will be on production support of capacity systems at LANL and planning for future TLCC architectures and integration.

In FY10, LANL procured the Hurricane-2 and Mapache systems, and deployed TOSS stack and tools on capacity systems.

Planned activities in FY11:

- Provide production support of LANL TLCC1 systems (Hurricane, Hurricane-2, Lobo, and Turquoise cluster)
- Move Lobo to SCC and integrate with Hurricane
- Install and integrate Hurricane-2 cluster in SCC
- Install and integrate Turquoise cluster in the Laboratory Data Communications Center (LDCC)
- Participate in evaluation and selection of TLCC2 vendor

Expected deliverables in FY11:

- Production support of LANL TLCC1 systems (Hurricane, Hurricane-2, Lobo, and Turquoise cluster)
- Lobo running in production mode and integrated with Hurricane
- Hurricane-2 cluster running in production mode in SCC
- Turquoise cluster running in production mode in LDCC
- Site preparation for TLCC2 SUs

Preliminary planned activities in FY12:

- Install and integrate TLCC2 systems
- Provide production support for all TLCC systems

WBS 1.5.4.2 ASC Capacity Systems (SNL)

The purpose of the ASC Capacity Systems project is to support the acquisition, delivery, and installation of new ASC capacity systems.

The project is supported by analysis of SNL's portfolio of application needs for capacity computing systems within the context of the higher integrated ASC platform strategy of capability, capacity, and advanced systems. Efforts include definition of requirements for TLCC system procurements and collaboration with *WBS 1.5.4.7 Common Computing Environment* with respect to a common software stack for new and existing capacity systems.

In FY10, SNL maintained and operated TLCC1 common capacity clusters with the CCE selected stack and tools in SNL's production computing environment. SNL supported the TLCC2 procurement process including requirements gathering, architecture requirements, vendor selection, and preparation of a Statement of Work.

Planned activities in FY11:

- Continue maintenance and operation of TLCC1 systems in SNL's production computing environment running the TOSS stack: Glory (2 SUs), Whitney (2 SUs), and Unity (2 SUs)

- Integrate CCE software stack and tools deployed on TLCC clusters
- Support tri-lab TLCC2 capacity systems procurement
- Support deployment of the tri-lab TLCC2 systems with appropriate Synthetic Workload and acceptance testing; integration into each lab's environment; and transition from initial Limited Availability to GA

Expected deliverables in FY11:

- TLCC1 systems operating in SNL's production computing environment running the TOSS stack
- Deployment of the tri-lab TLCC2 systems with appropriate Synthetic Workload and acceptance testing; integration into each lab's environment; and transition from initial Limited Availability to GA

Preliminary planned activities in FY12:

- Continue maintenance and operation of TLCC1 and TLCC2 systems in SNL's production computing environment running the TOSS stack

WBS 1.5.4.3: Advanced Systems

This level 4 product provides advanced architectures in response to programmatic, computing needs. The scope of this product includes strategic planning, research, development, procurement, testing, integration and deployment, as well as industrial and academic collaborations. Projects and technologies include strategic planning, performance modeling, benchmarking, and procurement and integration coordination. This product also provides market research, and the investigation of advanced architectural concepts and hardware (including node interconnects and machine area networks) via prototype development, deployment, and test bed activities. Also included in this product are cost-effective computers designed to achieve extreme speeds in addressing specific, stockpile-relevant issues through development of enhanced performance codes especially suited to run on the systems.

Advanced Systems Deliverables for FY11

- Sequoia Go/NoGo decision and subsequent work based on the decision
- Evaluation of Sequoia benchmarks on initial Sequoia hardware
- Second generation Sequoia R&D storage along with cloud computing into Hyperion environment
- Analyst support for weapons community users
- Introduction of fine-grained frequency modulation in an application library to assess performance loss versus power savings

WBS 1.5.4.3 BlueGene/P and BlueGene/Q Research and Development (LLNL)

The BlueGene/P and BlueGene/Q R&D project is a multi-year NNSA and Office of Science R&D partnership with IBM on advanced systems. It targets the development and demonstration of hardware and software technologies for 1-petaFLOPS and 10-petaFLOPS systems. The BlueGene/P hardware is based on an extension of the highly

successful BlueGene/L architecture with more cores per node, faster nodes, more memory, faster interconnects, and larger system scalability. The software approach to BlueGene/P is open-source collaborative development between IBM research, Linux Technology Center, the IBM Engineering and Technology Services Division, Argonne National Laboratory, and the ASC tri-labs. In FY08, a BlueGene/P system was delivered to Argonne. In FY09, a BlueGene/P system was delivered to LLNL as the first system (the Sequoia ID (Dawn)) in the Sequoia procurement. Follow-on BlueGene/Q system design targets a 20-petaFLOPS system at the end of the contract.

This project incorporates requirements from the DOE laboratories, especially Argonne and LLNL, to have input into design choices and system testing for microprocessors, node architectures, and interconnects. The DOE laboratories also provide critical input on software, ensuring appropriate capability and features for the design target.

In FY10, LLNL led the Argonne/LLNL partnership in its interactions with IBM on the topic of BlueGene/Q. We investigated compiler capabilities for utilization of floating point, thread level speculation, and transactional memory. We used the field programmable gate arrays (FPGA) simulation to gather performance data, laying the groundwork for the testing of benchmark code performance that will continue in FY11.

Planned activities in FY11:

- Continue technical interaction with IBM on hardware issues and software development
- Continue to investigate compiler capabilities for utilization of Quad floating point (QPX), thread level speculation (TLS), and transactional memory (TM)
- Utilize software and hardware simulation capabilities to gather estimates of performance for benchmark codes

Expected deliverables in FY11:

- Performance predictions for TLS and TM enhancements to benchmark codes
- Perform work associated with preparing second version of BlueGene/Q chip hardware and price projections for build Go/NoGo decision

Preliminary planned activities in FY12:

- Support Sequoia final phase 3 system build and early science period through technical interactions with IBM on hardware and software issues

WBS 1.5.4.3 Petascale Application Enablement (LLNL)

The Petascale Application Enablement project enables advanced application work to develop benchmarks for new platforms, such as Sequoia, and to adapt current codes to the expected new architectures. A primary target of this project is investigating ways to improve application thread and I/O performance for future many-core platforms. The project team efforts include both direct application work and benchmark development and testing.

In FY10, LLNL evaluated expected performance for Sequoia benchmark applications and kernels from key science applications on the Sequoia simulators. LLNL also continued investigations into threading performance issues with a particular focus on novel hardware features anticipated in Sequoia. Exploration of software transactional memory suggested that hardware support for this programming paradigm could benefit

several ASC applications. Performance improvements were realized on current architectures for key ASC applications, with a particular focus on integrated codes.

Planned activities in FY11:

- Continue vendor interactions with respect to Sequoia application performance requirements
- Continue science and weapons code testing on the Sequoia ID (Dawn) system and Sequoia simulators, especially enhancing single node thread performance
- Investigate opportunities for thread-parallel performance in production applications
- Identify mechanisms to enhance science and weapons code I/O performance

Expected deliverables in FY11:

- Evaluation of Sequoia benchmarks on initial Sequoia hardware

Preliminary planned activities in FY12:

- Initiate science and weapons code testing on the Sequoia system
- Continue with focus on code improvement opportunities

WBS 1.5.4.3 Sequoia (LLNL)

The Sequoia project will deploy a multi-petaFLOPS computer in early FY12 to be operated as an SSP user facility focused on supporting UQ and reduction in phenomenology (the elimination of code “knobs”). The primary missions of the machine will be (1) UQ for certification and model validation; and (2) weapons science investigations whose resolution is necessary for predictive simulation and, therefore, stockpile transformation. Sequoia will provide computational resources up to 12–24 times more capable than ASC Purple for UQ and up to 20–50 times more capable than BlueGene/L for weapons science investigations. Sequoia will bridge the gap between current terascale systems and later exascale systems that will become available within a decade. The Sequoia acquisition activity awarded build, delivery, development, and engineering contracts to IBM on January 8, 2009.

There are two major deliverables for Sequoia. The first deliverable (acquisition and delivery of an early environment—called the Sequoia ID (Dawn)) was accepted in CY09Q1 and dedicated on May 27, 2009. The second deliverable (acquisition and delivery of final Sequoia environment) is to be completed by end of 2011. These environments (both the Sequoia ID (Dawn) and Sequoia) will consist of a large compute platform, plus requisite federated switch networking infrastructure and parallel file system storage hardware (augmenting LLNL’s existing Lustre parallel file system deployments) to support compute platforms. Acquired switching infrastructure and storage hardware may also have high-speed hardware connectivity to servers and resources at LLNL outside of the Sequoia ID (Dawn) and Sequoia, including visualization engines, archival storage movers, BlueGene/L, Purple, and TLCC1 Linux clusters.

Supporting the final Sequoia platform deliverable is a D&E contract for hardware, software, and scalability items. The Sequoia acquisition RFP process requested technology roadmaps from responding vendors, with the intent that the Sequoia project would fund D&E efforts for the winning vendor to address the identifiable risks and issues for platform build and delivery. The D&E contract with IBM is intrinsic to overall

Sequoia risk mitigation strategies and helps to ensure the successful “productization” of the Sequoia components and systems.

In FY10, LLNL transitioned Sequoia ID (Dawn) to the SCF and provided a functional software environment for users. LLNL demonstrated the first pass Sequoia prototype. In addition, LLNL managed contractual milestones, and held monthly reviews for Sequoia planning and project management. LLNL continued to plan for Sequoia platform delivery.

Planned activities in FY11:

- Manage FY11 D&E contractual milestones
- Hold monthly review(s) for Sequoia planning and project management
- Prepare for Sequoia build Go/NoGo

Expected deliverables in FY11:

- Sequoia Go/NoGo decision
- Sequoia parts commit
- Sequoia parts and Phase 1 and 2 system builds

Preliminary planned activities in FY12:

- Sequoia Final Phase 3 system build in Q1 and Q2
- Final Sequoia System Acceptance in Q2
- Sequoia Early Science period in Q3
- Transition Sequoia to SCF in Q4
- Write Sequoia Operations Readiness (CD4) Package in Q4, pending approval in FY13 prior to production operations (GA)

WBS 1.5.4.3 Hyperion Test Bed (LLNL)

With the extremely demanding I/O requirements of petascale applications for Sequoia and the need for improved scientific data management capabilities, it is clearly apparent that emerging breakthrough technologies such as storage server virtualization and flash memory need to be tested in a large-scale environment such as Hyperion. The Hyperion Test Bed project will work with an expanded set of Hyperion vendor partners in the Hyperion Data Intensive Test Bed to evaluate innovative alternative storage and storage area network (SAN) architectures and push the achievable bandwidth I/O operations per second performance boundaries. In addition, it will work with the Hyperion Data Intensive Test bed to improve the fail-over and reliability of parallel file systems.

In FY10, LLNL added the following new capabilities: deployed new storage and networking hardware for the Hyperion Data Intensive Test Bed, including high performance flash memory, traditional disk storage systems, and next generation InfiniBand networking equipment, and Hyperion Memorandum of Understanding signed by all member institutions.

Planned activities in FY11:

- Evaluate the innovative storage and SAN architectures deployed in the Hyperion Data Intensive Test Bed

- Ongoing software and hardware scalability testing and evaluation by the Hyperion community

Expected deliverables in FY11:

- Second generation Sequoia R&D storage into Hyperion environment
- Cloud computing test bed into Hyperion environment
- Technology upgrade and refresh of Hyperion phase 1 SU

Preliminary planned activities in FY12:

- Support Sequoia storage and SAN deployment
- Hyperion moves to Livermore Valley Open Campus (LVOC) and available to broader community

WBS 1.5.4.3 Advanced Systems Technology Research and Development (SNL)

The Advanced Systems Technology R&D project will work in conjunction with the Institute for Advanced Architectures and Algorithms (IAA) to help overcome some of the bottlenecks that limit supercomputer scalability and performance through architectures and software research. The project's architecture efforts will focus on 1) advanced memory technologies, 2) high speed interconnects, and 3) power management techniques to reduce runtime power consumption of current and future platforms. The project will also explore software capabilities that will be essential for increasing the performance and scalability of applications beyond those provided by general-purpose operating systems, while providing the functionality necessary to deal with new compute node architectures, networks, parallel programming models, and applications—all at extreme scale.

This project will build upon researchers' expertise to increase application performance on future many-core processors, ease the transition of applications to alternative programming models, and provide robust system software support for scaling applications to a million-way parallelism. This will include work on enabling the user to interact easily with the data in order to perform analysis.

In FY10, SNL worked with industry, academia, and laboratory partners to develop advanced memory and high-speed interconnect subsystems to improve performance on legacy and new ASC applications. We held several workshops for ASC Integrated Code developers so they can position their codes for new technologies that are likely to appear in future ASC platforms.

Planned activities in FY11:

- Introduce fine-grained frequency modulation in an application library to assess performance loss versus power savings
- Continue research into runtime power management techniques, which will include exploring collaborative research partners
- Continue the development of advanced high-speed interconnect architectures capable of supporting exascale architectures
- Develop topology aware MPI enhancements for better utilization of next generation platforms such as Cielo and Sequoia.

Expected deliverables in FY11:

- Incorporation of fine-grained frequency modulation in a system library and publication of results

Preliminary planned activities in FY12:

- Incorporate frequency tuning into a more modern lightweight kernel (LWK), such as kitten; begin research with an external collaborator

WBS 1.5.4.4: System Software and Tools

This level 4 product provides the system software infrastructure, including the supporting operating system environments and the integrated tools, to enable the development, optimization, and efficient execution of application codes. The scope of this product includes planning, research, development, integration and initial deployment, continuing product support, and quality and reliability activities, as well as industrial and academic collaborations. Projects and technologies include system-level software addressing optimal delivery of system resources to end-users, such as schedulers, custom device drivers, resource allocation, optimized kernels, system management tools, compilers, debuggers, performance tuning tools, run-time libraries, math libraries, component frameworks, other emerging programming paradigms of importance to scientific code development and application performance analysis.

System Software and Tools Deliverables for FY11

- TOSS and Simple Linux Utility for Resource Management (SLURM) deployments on LLNL platforms
- SAP and outreach for Sequoia (Level 2 milestone)
- Applications development tools environment improvements and certification on Sequoia and TLCC2
- Application assessment for MPI usage to increase throughput
- Update of *Software Support and Development* capability strategy plan
- Science/code team application runs and regression testing completed to standup the Cielo system and move toward an operational state after it has moved to the SCF
- Science runs to standup and move the Conejo system into production
- Advanced data structure, KD-tree for structured 2D and 3D AMR and unstructured meshes, committed into Crestone project
- Advanced communication modules committed into Crestone project for the advanced parallel strategy
- Parallelization of the refactored ASC codes with the advanced parallel strategy implemented
- Parallel IO capability for visualization committed into the refactored ASC codes and associated visualization tools developed
- An assessment of the relevance of the motifs to ASC codes, and recommendations for their extension and generalization
- Tools for the analysis of code structure and dynamic behavior of applications on emerging architectures

- Proof of concept, application-level fault tolerance in an ASC physics code
- Improved performance metric and tracking capabilities delivered to code projects
- Investigate application-aware caching techniques and distributed file systems utilizing node-local cluster storage, both solid state (SSD) and rotating, to accelerate interactive DISC applications
- A prototype scalable / distributed / dynamic boot launch service for exascale architectures
- Implementation of compile- and run-time analysis tools supporting a selection of programming abstractions and models for emerging and extreme-scale computing systems
- Extensive support for stand-up of Cielo, including completion of acceptance criteria
- Demonstration of the ability to affect application through appropriate resource allocation based on historic, run-time, or user furnished process resource requirements in conjunction with current platform resource usage and availability
- Structural Simulation Toolkit (SST) v3.0 release

WBS 1.5.4.4 System Software Environment for Scalable Systems (LLNL)

The System Software Environment for Scalable Systems project provides system software components for all the major platforms at LLNL, research and planning for new systems and future environments, and collaborations with external sources such as the platform partners, especially IBM and Linux vendors. This project covers system software components needed to augment Linux and required proprietary operating systems that function in a manageable, secure, and scalable fashion needed for LLNL ASC platforms.

This project includes work on developing, modifying, and packaging the TOSS, and developing scalable system management tools to support the operating system and interconnect (for example, TOSS and InfiniBand monitoring tools), as well as the resource management environment (Moab and SLURM) to queue and schedule code runs across LLNL systems. LLNL uses TOSS on all of its Linux clusters. This project also funds approximately 60 percent of the manpower required to develop, deploy, and maintain TOSS. The funding LLNL receives for its portion of CCE TOSS funding accounts for 40 percent of the effort required to develop, deploy, and maintain TOSS. Therefore, TOSS activities and deliverables at LLNL are captured both here and in section 1.5.4.7 of this document.

In FY10, LLNL added the following new capabilities: deployed TOSS 1.3 (based on Red Hat Enterprise Linux (RHEL) 5.4), deployed OpenFabric Enterprise Distribution 1.4, released SLURM 2.1, and deployed Moab and SLURM to Sequoia ID (Dawn) system.

Planned activities in FY11:

- Continue ongoing development / support for TOSS software, including integration support for TLCC2 InfiniBand hardware
- Develop / deploy TOSS 2.0 (based on RHEL 6)
- Develop InfiniBand SAN support for Sequoia
- Continue ongoing development and support of Moab and SLURM

- Begin the preliminary work to port SLURM to Sequoia

Expected deliverables in FY11:

- TOSS 1.4 (based on RHEL 5.5) deployed
- TOSS 2.0 deployed
- SLURM 2.2 released to accompany the TOSS 1.4 release
- SLURM port to Sequoia in test

Preliminary planned activities in FY12:

- Continue ongoing TOSS software development and support
- Investigate future SAN technologies that support RDMA transports for more efficient storage traffic
- Continue ongoing development and support of Moab and SLURM
- Port SLURM to Sequoia system

WBS 1.5.4.4 Applications Development Environment and Performance Team (LLNL)

The Applications Development Environment and Performance Team (ADEPT) project provides the code development environment for all major LLNL platforms, supports user and code productivity, provides research and planning for new tools and future systems, and collaborates with external sources of code development tools such as platform partners, independent software vendors, and the open source community. The project works directly with code developers to apply tools to understand and to improve code performance and correctness. The project resolves bug and user trouble reports, including interactions with the software providers to fix problems.

The elements of the development environment covered by this project include, but are not limited to, compilers, debuggers, performance assessment tools and interfaces, memory tools, interfaces to the parallel environment, code analysis tools, and associated run time library work, with explicit focus on the development environment for large-scale parallel platforms.

Interactions between project members and code development teams ensure high performance use of existing systems and supports customer-based planning of future improvements to the environment. Similarly, long-term relationships with external partners, such as IBM, TotalView Technologies, the Krell Institute and Openworks, ensure that project members can resolve trouble reports quickly and avoid unnecessary duplication of existing capabilities.

In FY10, with the availability of the Sequoia simulators and the Sequoia ID (Dawn) system applications development tools environment, the ADEPT project included a high priority focused activity within ADEPT, entitled SAP, that will ensure that a wide range of ASC applications run effectively on Sequoia when it arrives, including exploiting the novel aspects of its hardware and software. It also includes ADEPT responsibility for the SAP Level 2 milestone. In addition, ADEPT has deployed production-quality versions of portions of the scalable code correctness suite, with a particular focus on the Stack Trace Analysis Tool. Further, ADEPT has begun planning for exascale systems, including an evaluation of the most critical emerging issues and strategies to address them.

Planned activities in FY11:

- Provide, maintain, and refine Purple, the Sequoia ID (Dawn), BlueGene/L, and TLCC2 code development environments
- Coordinate integrated design code scaling for Sequoia as part of SAP activities
- Explore novel threading strategies and performance as part of SAP activities
- Work with tri-lab partners to advance the common tri-lab environment for TLCC2
- Design and develop exascale development environment approaches
- Develop new techniques to improve robustness and performance of ASC codes
- Interact with the ASC code teams and vendors to improve software products

Expected deliverables in FY11:

- The FY11 SAP milestone, including characterizing the performance on the Sequoia ID (Dawn) system and Sequoia simulators and deployment plan and initial testing of Sequoia development environment
- Sequoia ID (Dawn) system applications development tools environment improvements
- Development of environment portion of exascale planning milestone
- Deployment of a full production version of highly scalable code correctness tool suite
- TLCC Tripod applications development tools environment certification for new TOSS releases

Preliminary planned activities in FY12:

- Continue code development environment support on all LLNL ASC platforms
- Identify and develop refinements of the code development environment for existing and future capacity and capability systems
- Continue to explore nested node concurrency programming model interfaces and performance
- Continue to support users and to interact with vendors to serve user needs
- Participate in validation and acceptance of the Sequoia system

WBS 1.5.4.4 Software Support (LANL)

The Software Support project works to establish a strong development and analysis tool capability for current and next generation HPC platforms, including parallel capabilities. It is focused on working with the HPC tool community and vendors to identify, plan, and integrate tools into production environments and establish a solid support structure. The project supports the incremental improvement of tools driven by the ASC strategic plan. The plan includes cross-laboratory partnerships and external collaborations that focus on performance tools required for programming model support.

In FY10, LANL added a contract for Valgrind support and added support/training contract for TAU performance analysis suite. Added capabilities include tool strategic plan development based on current and next generation platform planning, integration

with HPC community tool development efforts, and tool development and production integration capability.

Planned activities in FY11:

- Continue building an MPI support capability by engaging the community support model; focus will be on OpenMPI for development needs
- Assess application MPI usage by platform architecture and recommend parameter additions or code enhancements
- Continue performance analysis support capability; focus will be on OI SS integration in application analysis and new analysis needs
- Increase memory analysis and MPI performance/correctness tool support
- Continue building a stronger debugger support capability; work with TotalView on tool deployments and enhancements; work with LLNL debugger capability through STAT (Static Analysis Tool) and subset debugging
- Continue involvement with DOE NNSA and Office of Science laboratories in exascale tool planning and development programs
- Coordinate software tools implementation for the Cielo deployment
- Gather ACES requirements for unclassified computer supported collaboration expected in FY11 and beyond, and report to ACES management for further planning
- Update of software support and development capability strategy plan

Expected deliverables in FY11:

- Further deployment of MPI analysis tools
- Application assessment for MPI usage to increase throughput
- Interaction with community and vendors to increase MPI scale capability

Preliminary planned activities in FY12:

- Continue building an MPI support capability by engaging the community support model
- Continue performance analysis support capability
- Continue building a stronger debugger support capability
- Continue involvement with DOE NNSA and Office of Science laboratories in tool planning and development programs

WBS 1.5.4.4 Applications Readiness (LANL)

The Applications Readiness project addresses issues with an application's production-run readiness on current and incoming computing systems at LANL. Working with subsystem teams such as systems management, file systems and I/O, archive, and tools, the Applications Readiness team identifies causes of unexpected behavior and deploys fixes in production. The project goal is that system users are able to make productive use of the systems with their applications to solve their problems.

The project provides production problem solving (create small problem reproducers, identify cause, work with the relevant vertical(s) to find a solution, and verify the deployed solution), periodic stress testing/regression of production machines, new

software version regression testing, system configuration verification and software stack deployment with real user applications and metrics, and analysis/profiling.

In FY10, LANL added the following capabilities: expanded the skill set of the team to include world-leading skills in hybrid architecture based programming techniques. These skills were developed by working with multiple application teams to migrate codes to Roadrunner and stabilize them in production. Our efforts also added new capabilities, in the form of higher throughput, in the Crestone Project by refactoring I/O subsystem and optimizing use of collectives.

Planned activities in FY11:

- Take on the hardest and most elusive problems, or newly identified problems, with existing production machines
- Stand up Cielo, Conejo, and other systems
- Continue to assist users with migration towards the use of hybrid programming to exploit the cell processors in the Roadrunner phase 3 system and monitor other significant developments in the hybrid computing space
- Perform special projects, including library code port to Roadrunner, parallel log structured file system (PLFS) testing, and SPaSM data intense
- Assist system management personnel with problem investigation and resolution

Expected deliverables in FY11:

- Science/code team application runs and regression testing completed to standup the Cielo system and move toward an operational state after it has moved to the SCF
- Science runs to stand up and move the Conejo system into production

Preliminary planned activities in FY12:

- Continue to take on the hardest and most elusive problems, or newly identified problems, with existing production machines
- Continue to assist users with migration towards the use of hybrid programming to exploit the cell processors in the Roadrunner Phase 3 system
- Assist system management personnel with problem investigation and resolution

WBS 1.5.4.4 Productivity Project (LANL)

The Productivity Project provides direct support to LANL ASC code projects for their productivity, capabilities, and performance on current and future ASC machines. This project proposes, implements, and tests new approaches to improve the productivity of ASC code projects on current and future computer architectures. These approaches include resilience through optimization procedure and scheduling, new mesh data structures, new strategies for AMR, performance improvement, data locality, compressed data structures for materials, and parallel strategies on advanced computer architectures.

This project provides a bridge between advanced computer science and ASC code projects for efficient usage of current and future computer hardware. The staff members in this project are able to identify main bottlenecks of current ASC code projects for better performance in calculation speed, memory usage, and communications, to develop strategies and plans for remedies, and to implement changes within ASC code projects. The project will focus on mesh data structures, compressed material data

structures, advanced parallel strategies, advanced mesh partition and load balance, and users' data structures for parallel I/O.

In FY10, LANL added the following new capabilities: developed the flattened mesh data structures within Crestone project, demonstrated the advantages of flattened mesh data structures over the existing hierarchy structures within Crestone project, developed the weight solution scheme with the library AMTTI, collected application data and validated models/tools against real data, developed a framework of compressed data structure for materials within Crestone project, demonstrated its capability through hydro calculations, and evaluated a new framework of AMR.

Planned activities in FY11:

- Design an advanced parallel strategy in refactored ASC codes
- Develop parallel KD-tree data structures in refactored ASC codes
- Put the new data structures in wide use within applications of refactored ASC codes
- Implement the advanced parallel strategy in refactored ASC codes for much better performance
- Develop advanced parallel I/O capability for visualization in refactored ASC codes
- Develop hydro capability within the new framework of AMR

Expected deliverables in FY11:

- The advanced data structure, KD-tree for structured 2D and 3D AMR meshes committed into Crestone project
- KD-tree for unstructured meshes committed into Crestone project for mapping
- Advanced communication modules committed into Crestone project for the advanced parallel strategy
- Parallelization of the refactored ASC codes with the advanced parallel strategy implemented
- Parallel I/O capability for visualization committed into the refactored ASC codes
- Tools to visualize the visualization dumps developed
- Published research papers about the new AMR strategy

Preliminary planned activities in FY12:

- Run refactored ASC codes on full or almost full most recent ASC computers
- Improve parallel performance of the refactored ASC codes
- Improve the AMR technology used in current ASC codes
- Develop advanced AMR technology

WBS 1.5.4.4 Programming Models for the Next-Generation Scientific-Computing Environment (LANL)

The goal of this project is to study new hardware directions, programming abstractions and models, the software development tool chain, and run-time systems for scientific-computing environments with the goal of developing a set of techniques that will assist in the development of the next-generation of applications.

Motivated by the on-going revolution in computer architectures featuring the emergence of homogeneous and heterogeneous many-core processors, this project will explore and define programming models and abstractions, together with the supporting compile- and run-time infrastructure needed to provide the scientific-computing community with the necessary methods and tools to address the challenges of developing and executing applications on extreme-scale, highly-concurrent systems. Work will be guided by the needs of ASC IC and workloads, and common computation and communication patterns.

This project builds on the work of the Advanced Architectures and Usable Supercomputing project that is concluding in FY10. An important goal is to develop collaboration with SNL and LLNL on common challenges for the scientific-computing environment that are faced by each of the tri-labs in preparation for extreme-scale computing systems.

This is a new project for FY11.

Planned activities in FY11:

- Investigate programming abstractions for emerging architectures, with applications to multi-physics and data-intensive applications
- Develop techniques for the analysis of code structure and dynamic behavior of realistic (multi-physics) applications, and the identification of target code elements for acceleration on emerging architectures
- Develop tools and techniques for the analysis of data movement and performance of scientific applications
- Establish collaborations with SNL and LLNL to address issues in scientific computing at extreme scale
- Provide support to HPC Division in the deployment of the tri-lab production system Cielo, focusing on our areas of expertise: optimization, scalability, and networking

Expected deliverables in FY11:

- Example implementations of programming abstractions and models for scientific computing on emerging architectures and extreme-scale systems
- Implementation of compile- and run-time analysis tools supporting a selection of programming abstractions and models for emerging and extreme-scale computing systems
- Tools for the analysis of code structure and dynamic behavior of realistic applications
- Support to HPC Division for stand-up of Cielo, including completion of acceptance criteria

Planned activities in FY12:

- Refine programming abstractions, run-time systems, and supporting tools in the light of maturing understanding of architectural and application needs

WBS 1.5.4.4 Systems Research (LANL)

This project's objective is to develop exascale resource management capability in support of ASC systems and applications.

With future exascale systems in mind, we need to develop ways to manage the resources associated with billions of cores and to efficiently provide services at this scale. To meet this challenge, we will investigate scalable-dynamic-distributed algorithms to support such services as booting, system monitoring, job launch, job management, authority management, and services discovery. The design will take into account the specific needs that are key to success with exascale: fault tolerance, power conservation, and sheer scale. We will also investigate what some of the key factors are that limit ASC's application scaling to exascale and the system resources needed for these applications. We propose to develop, test, and compare different algorithms necessary to support the services that are integral to reliable application execution on exascale class machines. We will also investigate methods of validating algorithms for exascale systems. Other related work would include:

- Investigate/evaluate cutting edge hardware and networks for applicability to exascale systems
- Systems anomaly determination and investigation in support of near term and future production systems

Planned activities in FY11:

- Investigate scalable-dynamic-distributed algorithms
- Develop a boot launch service, which will be developed and tested on local clusters

Expected deliverables in FY11:

- A prototype scalable/distributed/dynamic boot launch service for exascale architectures to include:
 - Efficient self-identification (DHCP replacement)
 - File Transfer—scalable delivery of boot environment and executables
 - Role-based booting—dynamic configuration of running services
 - Boot verification—verify booted systems and services

Planned activities for FY12:

- Extend scalable-dynamic-distributed exascale resource management research into other services areas, such as system monitoring and dynamic service initiation and discovery

WBS 1.5.4.4 Data Intensive Supercomputing Project (LANL)

Data-Intensive SuperComputing (DISC) aims to handle the “data deluge,” our generation of massive amounts of simulation data that threatens to outstrip our ability to access, organize, and manage it. Extreme-scale ASC databases—petabytes or larger—already exist and are growing in number and size. The databases are so large that data cannot be moved to HPC platforms for interactive analysis, as the transfer could take weeks or longer. Any non-trivial operation on the database must be done in place: the resulting DISC architecture is therefore organized around many thousands of disks in a networked cluster environment, in contrast to a conventional in-memory computation optimized, HPC platform.

What is desired is a combination of a transactional database and a parallel compute engine—the DISC system. Driven by this imperative, our long-term objective is to build a system that, within reasonable limits, can manage DISC performance optimization

tasks automatically, be fault-tolerant, provide a set of high-level serial and parallel programming primitives for carrying out complex queries and computations on the stored data, and provide fast enough execution to enable interactive, iterative discovery. Current DISC approaches range from the database-driven (SQL), to parallel computation-driven (Hadoop, Google's MapReduce), neither of which alone can meet the needs of the target community. The database approach does not map well to science data types, lacks required features (provenance, uncertainty), and cannot handle many standard scientific analysis operations. The stripped-down MapReduce approach, while scalable, needs substantial enhancement for science applications. We will build on these technologies integrating the best of both approaches.

Planned activities in FY11:

- Study emerging solutions in the area of DISC and evaluate their usefulness to ASC programmatic needs
- Prototype data intensive software aligned with emerging storage and supercomputing platforms to advance petascale data infrastructure

Expected deliverables in FY11:

- Partner with university collaborators and industry to explore the use of new DISC approaches on massive ASC data sets
- Investigate both application-aware caching techniques and distributed file systems utilizing node-local cluster storage, both solid state (SSD) and rotating, to accelerate interactive DISC applications

Planned activities in FY12:

- Develop a data intensive infrastructure that meets ASC programmatic needs by leveraging existing and custom data intensive software

WBS 1.5.4.4 Application Enablement for Extreme-Scale Computing (LANL)

This project will develop patterns, strategies, and abstractions for the implementation and optimization of scientific applications and algorithms on emerging hardware architectures, and to enable ASC integrated code projects to take the necessary steps to prepare for the next-generation of computer systems.

Due to the upheaval caused by emerging homogeneous and heterogeneous many-core processor architectures, the scientific-computing community is faced with both redesigning and re-thinking the way we develop applications for high-performance, large-scale systems. The project will track the development of emerging hardware architectures, and study both computational and data movement patterns represented by “mini apps,” including a subset of the thirteen “motifs” popularized by researchers at Berkeley, and exemplar applications developed in conjunction with ASC IC code developers. Using a co-design approach, we will explore issues of software abstraction, which system features can usefully be exposed at the application level, and the potential for advanced hardware support for application data movement and resiliency. We will also work with the ASC teams to enhance existing performance tracking tools.

This project builds on the work of the Advanced Architectures and Usable Supercomputing, and the Code Performance and Throughput projects that are concluding in FY10.

This is a new project for FY11.

Planned activities in FY11:

- Work with ASC code teams to develop exemplar applications to add to our library of mini-apps for further study
- Study mini-apps, their effective implementation on emerging architectures, and work with ASC code teams to identify useful abstractions for multi-physics applications on extreme-scale computing systems
- Work with ASC code teams to develop an autonomous, flash-based check-pointing system
- Provide performance metric and tracking tools for ASC integrated codes on current and emerging platforms
- Establish collaborations with SNL and LLNL to address in scientific computing at extreme scale
- Provide support to the HPC Division in the deployment of the tri-lab production system Cielo, focusing on areas of expertise—application-level understanding of performance and optimization

Expected deliverables in FY11:

- A library of mini apps representative of important issues for ASC applications
- Implementations of a number of mini apps on several emerging architectures
- Useful abstractions for the implementation of mini apps
- Demonstration of an autonomous, flash-based check-pointing system
- Improved performance metric and tracking capabilities delivered to code projects
- Support to HPC Division for stand-up of Cielo, including completion of acceptance criteria

Planned activities in FY12:

- Continue tracking the evolution of emerging architectures and the associated impact on exemplar applications
- Explore more mini apps and more architectures; extend lessons learned from analysis of mini apps to realistic applications
- Explore application of intelligent hardware to memory and data movement optimization in ASC codes

WBS 1.5.4.4 Software and Tools for Scalability and Reliability Performance (SNL)

The Software and Tools for Scalability and Reliability Performance project supports software R&D to address scalability and reliability of future computational systems. Since Cielo is in our immediate future, we will devote resources to ensuring its success.

Red Storm system software R&D and experiences will be applied to the deployment efforts of Cielo as it begins full production capability. Recent (FY09 and FY10, as well as FY11) system software, I/O, visualization, and simulation R&D efforts will be marshaled and finalized to assess the impact on extending the life and performance of ASC codes, without significant application changes. Performance modeling activities will evolve from the Cielo 6X testing to developing an analytical method for predicting

performance of future platforms. This work will become synergistic with the system simulation R&D as new processor models, mini applications, and failure modeling techniques are developed and integrated with the SST. Additional reliability and resilience R&D will utilize data from existing platforms to improve on statistical methods for associating failures with indirect causes and to associate job failures with system observables (for example, logs and network counters).

More advanced R&D efforts will focus on system software features to facilitate a transition to new programming models and improve resiliency through techniques such as failure prediction. We will employ various judicious and lightweight resource monitoring and data collection techniques to provide resource utilization feedback directly to applications as well as input to an intelligent modeling system that can subsequently alert applications. For instance, if a certain ensemble of information comprising an “observation” on a compute resource fits a model that is associated with high probability of failure, an event informing the application(s) running on that resource of the impending failure could be notified to take appropriate action, whether it be a live process migration or restarting a checkpoint with a replacement resource. Appropriate application programming interfaces (APIs) must be developed for these activities.

In FY10, due to the highly scalable nature of the Cray software stack, SNL was able to focus their risk mitigation efforts on a smaller subset of the Cielo system software than originally anticipated. The risk areas were identified as the networking software and the I/O capability. Performance analysis work focused on the 6X testing for Cielo. We implemented and demonstrated a systems approach prototype for analyzing and sharing information about resource and application state and for invoking live process migration via the mechanism of virtual machines to facilitate rebalancing process-to-resource mapping for a testbed scale MPI application. We have also improved scalability of our monitoring and analysis system as well as added further analyses.

Planned activities in FY11:

- Support evolutionary system software for current HPC hardware
- Address scalability issues in Cielo
- Develop performance prediction analytical method, mini applications, and processor models
- Enhance simulator of failures in HPC systems
- Enhance analysis of failures to root causes through statistical methodologies and using system state data
- Enhance the architecture of our monitoring and analysis system to enable run-time modeling and anomaly / degradation detection at large scale (> 10,000 computational units)

Expected deliverables in FY11:

- Metrics on application porting and performance for Cielo
- Performance prediction methodology
- Integration of failure models in SST
- Demonstration of high-fidelity monitoring with distributed data collection, run-time analysis, response trigger generation, and response at large scale (>10,000 computational units)

Preliminary planned activities in FY12:

- Integrate data collection tools with other system tools on target platform resources
- Participate in prototype deployment of system software tools and enable applications to measure performance improvements
- Expand monitoring and analysis to include system infrastructure components (for example, networks and storage)
- Incorporate root cause analysis into the system monitoring, analysis, and response framework

WBS 1.5.4.4 System Simulation and Computer Science (SNL)

Given the extreme cost of deploying a capability machine, as well as the high cost of developing the complex multi-physics codes that run on them, it is important to use a systematic design and evaluation approach that permits decisions to be informed by predictions of application performance on these machines. Such an approach can be used in architecture design, algorithm design, and platform procurement. The System Simulation and Computer Science project, in collaboration with the IAA, will continue the development of a multi-scale simulation capability within the context of the SST.

The SST consists of a core set of components that enable parallel discrete-event simulation; high-fidelity networking, memory, and processor components; and coarse-grained simulation components that capture essential elements of machine performance with low computational cost. Future HPC systems and the applications designed to utilize them are impacted by a variety of considerations, including scalability of applications, ease-of-programming, memory and network latencies becoming more imbalanced relative to computation rates, silent data corruption and its propagation, frequency of interrupts, power consumption, and overall machine cost. SST is designed to allow each of these parameters to be explored, allowing the consideration of a broad space of potential architectural and algorithmic designs. The goal is for the SST components to be extended and enhanced by a community of simulator developers including academic, industrial, and government partners. An even larger community is expected to be the users of SST, including algorithm developers, architecture designers, and procurement team members.

In FY10, SNL received copyright permission for SST and released the code under an open-source license. We have continued interactions with the Atomic Weapons Establishment simulation experts and have made our code available to them. The MPI tracing component of SST, DUMPI, was selected by Cray as the library they used for collecting data for Cielo. We worked closely with Cray to ensure it met their needs.

Planned activities in FY11:

- Improve SST's scaling, documentation, and usability
- Provide facilities for modeling faults and cost
- Integrate visualization and design space exploration tools
- Provide new models for memory, network, and processors
- Validate simulator using Cielo data and Mantevo applications

Expected deliverables in FY11:

- SST V3.0 release

Preliminary planned activities in FY12:

- Integrate HDL simulation capabilities into SST core
- Integrate FPGA co-simulation capabilities into SST core
- Prepare for release of SST V4.0

WBS 1.5.4.5: Input/Output, Storage Systems, and Networking

This level 4 product provides I/O (data transfer) storage infrastructure in balance with all platforms and consistent with integrated system architecture plans. The procurement of all supporting subsystems, data transfer, storage systems, and infrastructures occurs through this product. The scope of this product includes planning, research, development, procurement, hardware maintenance, integration and deployment, continuing product support, quality and reliability activities, as well as industrial and academic collaborations. Projects and technologies include high-performance parallel file systems, hierarchical storage management systems, storage-area-networks, network-attached storage (NAS), and high-performance storage system (HPSS) or future hierarchical storage management system disks, tape, robotics, servers, and media. This product also includes relevant prototype deployment and test bed activities. Projects and technologies in the advanced networking and interconnect areas include networking and interconnect architectures, emerging networking hardware technologies and communication protocols, network performance/security monitoring/analysis tools, and high performance encryption and security technologies.

Input/Output, Storage Systems, and Networking Deliverables for FY11

- Sequoia file system procurement
- Networking preparation for Sequoia
- Lustre deployments on LLNL platforms
- Release HPSS 7.4, 7.5, and 8.0 as appropriate
- Archival storage hardware and software procured and deployed for ASC platforms
- Parallel Storage Interface (PSI) 2011 release
- Deployment of PLFS on Cielo
- Production-level release and deployment of GPGPU RAID on SNL hardware
- Prototype I/O services to perform *in-situ* analysis for CTH and/or Charon

WBS 1.5.4.5 Archive Storage (LLNL)

The Archival Storage project provides end-to-end long-term, high-performance, archival storage services to ASC customers. This includes a collaborative software development effort between the tri-labs, Oak Ridge National Laboratory, Lawrence Berkeley National Laboratory, and IBM, as well as deployment and support of archival storage software and interfaces for tri-lab ASC customers on unclassified and classified networks. It also includes the selection, procurement, deployment, support, and maintenance of archival storage hardware and storage media and the ongoing technology refresh and data

stewardship. Archival storage system software (currently, HPSS) provides scalable, parallel archival storage interfaces and services to customers running at the tri-labs. HPSS distributes data across a configurable amount of storage units and removes other limits to scaling including number of files, directories, and concurrent users.

A world-class array of storage hardware is integrated beneath HPSS, supplying the performance necessary to offload ASC platforms, thereby increasing computation. This includes disk arrays, tape subsystems, mover nodes, storage-area-networks, networks, robotics and petabytes of media. Together, this hardware and software supports high-speed parallel transfer rates (that are currently in excess of 5 GB/sec. at LLNL) into a virtually unlimited data store, at a current capacity of over 17 PB in a single name space.

In FY10, HPSS R7.2 and R7.3 were developed and released and HPSS R8.1 design was finalized and development began. HPSS R7.3 was deployed to LLNL production environments featuring “aggregate-aware” repack enhancements and user defined attributes enabling custom metadata associations for HPSS files. Deployment teams also selected, procured, and deployed new disk cache capability (960 TB and 10 GB/sec.), as well as tape library expansions to increase the cartridge capacity of LLNL production archives from 130,000 tape slots to 184,000 tape slots. Obsolete hardware (including a set of nine robotic tape silos) was decommissioned throughout the year.

Planned activities in FY11:

- Develop HPSS R7.4, a minor release including conversion to a 64-bit architecture, support for IPv6, and new capability to repack existing tape-resident small files into aggregates
- Develop HPSS R7.5, a minor release with Redundant Array of Independent Tapes (RAIT), which will provide higher reliability more cost-effectively than multiple copies on tape
- Develop HPSS R8.1, featuring an architecture utilizing distributed core servers and partitioned metadata to meet the extreme scalability requirements of petascale computing
- Deploy programmatic quota system for LLNL archive users
- Deploy a later release of HPSS (R7.3.x or R7.4.x) in production environments
- Provide ongoing support of currently deployed archival storage systems, including selection, deployment, support and maintenance of all archival storage hardware and media, customer and interface support, ongoing tech refresh, and data stewardship of LLNL archives

Expected deliverables in FY11:

- Release of HPSS R7.4 and R7.5
- Coding and testing of HPSS R8.1
- Production deployment of a later release of HPSS (R7.3.x or R7.4.x)
- Archival storage hardware and software procured and deployed to support expansion of ASC platforms

Preliminary planned activities in FY12:

- Develop HPSS R8.1 through code review, system test, integration test phases, and beta release
- Plan for production deployment of HPSS R7.5

- Begin beta testing HPSS R8.1 in LLNL pre-production environment

WBS 1.5.4.5 Parallel and Network File Systems (LLNL)

This project provides for the development, testing (feature, capability, performance, and acceptance), procurement, integration, and ongoing support of various file system technologies and interfaces necessary for the efficient and effective use of ASC high-performance platforms. Included are the continuing development and support of Lustre as a fully-featured file system for the range of ASC capability and capacity platforms, the deployment and support of global parallel file system (GPFS), the deployment and support of ubiquitous NAS services for home, project, and scratch space, and the I/O support of various programming interfaces for parallel I/O.

This project deploys and supports Lustre and GPFS file systems for ASC platforms as well as high-availability NAS file systems for home and project space, and scratch space for serial capacity clusters. It actively works with Oracle Corporation and the Lustre development community to add Lustre file system scalability and reliability enhancements required by TLCC2 and Sequoia platforms. The file system up through the programming interfaces are supported to help developers of applications use parallel I/O effectively.

In FY10, LLNL co-developed and deployed Lustre Release 1.8, which focused on availability enhancements including Version-Based Recovery and per-user quotas. Lustre deployment teams procured, deployed, and operated a new multi-petabyte file system and expanded the file system infrastructure supporting the Sequoia ID (Dawn). A new data analysis NAS system was deployed for the Graph machine and home directory servers were fully upgraded. The parallel file system risk mitigation study was completed resulting in the transition of the lscratchc file system to GPFS.

Planned activities in FY11:

- Deploy ZFS-based Lustre prototype file system into production
- Provide Lustre file system performance and scalability enhancements in support of Lustre 1.8.4 and initial 2.x releases
- Work closely with application teams and Sequoia developers on implantation and performance of the Lustre client on Sequoia hardware
- Evaluate, specify, and complete procurement of a file system in support of Sequoia
- Expand production Lustre file systems in preparation for TLCC2 computational increases
- Upgrade the open computing facility (OCF) data analysis file system and deploy new project NAS servers on the OCF
- Maintain and support GPFS, NAS, and Lustre parallel file systems, including middleware and higher-level I/O libraries for users

Expected deliverables in FY11:

- Deployment of prototype ZFS-based Lustre file system
- Development and deployment of Lustre 1.8.4 and 2.x
- Replacement of data analysis and project NAS servers in the OCF
- Establishment of Sequoia file system procurement contract

Preliminary planned activities in FY12:

- Deploy Sequoia file system into full production
- Continue file system expansion based on TLCC2 computational needs
- Deployment of Lustre 2.x release into production in OCF and SCF environments

WBS 1.5.4.5 Networking and Test Beds (LLNL)

The Networking and Test Beds project provides research, performance testing, capability testing, and analysis for the file system, network, and interconnect subsystems in support of current and future systems and environments. This work relies heavily on an adequately provisioned test bed, skilled staff, and collaborations with vendors.

This project will test various hardware and software components to quantify the features, performance, reliability, security, and interoperability of the products and broader technology base. The information acquired as a result of this project will be used to help determine an integrated architecture and resultant procurements for these subsystems.

In FY10, host adaptor problems were resolved to improve stability of the network. Beta test of DataCenter Ethernet was initiated to verify functionality of newer switches. Throughput testing of next generation firewalls was successfully performed.

Planned activities in FY11:

- Perform research and testing for technologies and products for interconnects, local area networks (LANs), WANs, and National Security Agency (NSA) Type 1 encryptors, file system servers, clients and disks, with special focus on emerging DataCenter Ethernet switches, additional features in InfiniBand, other interconnect and 10GigE-related technologies, and 10GigE NSA Type 1 encryptors in support of future ASC petaFLOPS systems
- Study developing load balancing and multipath routing in support of congestion avoidance for previously listed networks
- Apply testing results to optimize the functionality, performance, reliability, manageability, and security of the I/O services supporting these computing systems

Expected deliverables in FY11:

- Switching infrastructure upgrade for future Sequoia network augmentation
- Preparation for procurement and initial installation of Sequoia network

Preliminary planned activities in FY12:

- Continue to leverage tri-lab activities in I/O-related hardware and software, and seek to improve the reliability, performance, and manageability of the I/O subsystems in production
- Research and test to determine which technologies and products should be considered for insertion into production to meet the growing I/O performance and capacity requirements
- Continue network improvements and track emerging network technology, especially related to congestion

WBS 1.5.4.5 File Systems and Input/Output Project (LANL)

The File Systems and I/O Project provides end-to-end, high-performance networking and scalable I/O infrastructure for the ASC program. It also delivers high bandwidth, low-latency interconnect technologies for the ASC compute platforms. The ASC program requires system and SAN bandwidths at over 500 GB/sec., global file system I/O rates beyond 500 GB/sec., and latencies in the 1 microsecond range. All this performance must be provided in an integrated, usable, reliable, and secure way. Data transfer and storage bottlenecks are still a critical concern for current and next-generation, HPC environments. Successfully meeting the ASC programmatic milestones requires carefully balanced environments in which the I/O infrastructure scales proportionally with increased ASC platform capabilities and application data needs.

This project is a coordination point for planning of all online storage, network, and data movement activities within the ASC program at LANL. These capabilities include online file systems such as the network file system (NFS) complex and enterprise-wide supercomputer file systems, GPFS development, deployment and management, scalable I/O middleware development and support, interconnect technology development and deployment, and SAN development and deployment.

In FY10, this project extended Roadrunner's capabilities by providing a fully functional data path to storage that includes the I/O nodes, PaScalBB infrastructure, and storage. ASC applications now have the capability to utilize PLFS. Acceptance of the Cielo file system hardware and early testing started in support of deployment of the ASC Cielo machine at LANL and the PaScalBB network architecture in the secure, successfully expanded to 12 lanes in preparation for Cielo.

Planned activities in FY11:

- Evaluate quality of service performance guarantees for file systems
- Prepare and test I/O infrastructure for the Cielo system in the secure

Expected deliverables in FY11:

- Collaboration with production team on deployment of PLFS on Cielo

Possible deliverables in FY12:

- Support deployment of the PLFS with metadata hashing

WBS 1.5.4.5 Archival Storage Design and Development (LANL)

The Archival Storage Design and Development project includes services for HPSS and PSI software development by LANL for the purpose of supporting ASC customers from LANL, LLNL, and SNL. These services include collecting user requirements for changes and upgrades to HPSS and PSI, developing plans for implementing user requirements into the codes performing the design and development work for upgrading the codes, and providing second-level support for the archive storage deployment team. The project works with the consulting office and archive storage deployment team to troubleshoot problems experienced with storing and retrieving data from the archive.

The HPSS portion collaborates with tri-lab developers for implementing solutions that meet ASC requirements for all three labs. The PSI portion collaborates with LANL colleagues on user interface issues and ensures that PSI functions with each new release of HPSS.

HPSS is software that manages petabytes of data on disk and robotic tape libraries. HPSS provides highly flexible and scalable hierarchical storage management that keeps recently used data on disk and less recently used data on tape. HPSS uses cluster, LAN, and/or SAN technology to aggregate the capacity and performance of many computers, disks, and tape drives into a single virtual file system of exceptional size and versatility. This approach enables HPSS to meet otherwise unachievable demands of total storage capacity, file sizes, data rates, and number of objects stored.

HPSS provides a variety of user and file-system interfaces ranging from the ubiquitous VFS, FTP, SAMBA, and NFS to higher PFTP, client API, local file mover, and third party SAN (SAN3P). HPSS also provides hierarchical storage management services for IBM GPFS.

In FY10, LANL added the following new capabilities: PSI 2010 release, redefined release strategy for HPSS, refined requirement documents for HPSS 8.0 release, and 7.3 release of HPSS.

Planned activities in FY11:

- Finalize requirements for 2011 release of PSI
- Design, develop, and test 2011 release of PSI
- Develop requirements for 2012 release of PSI
- Release HPSS 7.4
- Release HPSS 7.5
- Final design requirements for release 8.0 of HPSS
- Design, develop, and test HPSS 8.0
- Provide short-term functional updates to the existing HPSS code base
- Provide second-level support for archival storage deployment team

Expected deliverables in FY11:

- HPSS 8.0 release
- PSI 2011 release
- Release HPSS 7.4
- Release HPSS 7.5

Preliminary planned activities in FY12:

- Continue Level 2 support for the production archive
- Initiate next-generation archive planning

WBS 1.5.4.5 Archival Storage (SNL)

The Archival Storage project represents SNL's participation in the DOE HPSS Consortium development project. HPSS provides the archival storage solution for ASC systems and is in direct alignment with ACES.

SNL's role in the HPSS project is to collaborate with tri-lab developers to design, implement, and test solutions that meet ASC requirements for all three labs.

In FY10, SNL officially released HPSS version 7.3 and began development of versions 7.4 and 7.5, which includes support for RAIT. SNL also began design and prototyping of version 8.1, which includes support for multi-core servers for HPSS. Initial results on the 8.1 prototype demonstrated a file create rate of 1000 files/sec./server and appear to scale linearly up to eight servers.

Planned activities in FY11:

- Provide support for PFTP, HACL, HPSSD
- Continue design, development, testing for 8.1, multi-core-server, RAIT

Expected deliverables in FY11:

- Coding support for Version 8.1 of HPSS
- Requirements definition for next HPSS upgrade

Preliminary planned activities in FY12:

- Extend capabilities of HPSS version 8 by adding support for content validation, user-defined attributes, high-availability heartbeat and failover, and multilevel security

WBS 1.5.4.5 Scalable Input/Output Research (SNL)

The Scalable I/O Research project will drive I/O system enhancements for the next generation of extreme-scale systems through R&D in parallel file systems and I/O libraries.

This project will explore the use of hardware accelerators to improve bandwidth and reduce latency of I/O operations; investigate ways to exploit emerging storage architectures, such as non-volatile storage devices and data-warehouse appliances to improve I/O; and research ways to use available system resources to provide file-system caching and advanced in-transit data processing for I/O functions and high-level I/O libraries.

It is expected that most of the FY10 R&D of parallel file systems and I/O libraries will yield candidate technologies for future exascale systems. These technologies, coupled with advances in network transport services, such as remote direct memory access and service guarantees, will create new opportunities to push state-of-the-art in I/O.

In FY10, SNL made significant contributions to ASC as well as the broader HPC community. The research project to investigate GPGPU RAID received acceptance from the Linux community, and SNL began plans to transition to a production product, with testbed installations planned at SNL. The success of the RAID project also led to a collaboration with LSI to design and develop next-generation RAID controllers. The scalable I/O services work began collaborative efforts with Oak Ridge National Laboratory to develop in-transit I/O services to support fusion codes as well as the CTH shock physics code. The SICAIDA project completed and published a study evaluating the effectiveness of data-warehouse appliances for scientific analysis, and it began work to develop *in-situ* analysis services for the S3D combustion code.

Planned activities in FY11:

- Support I/O and file system issues for Cielo, TLCC, and testbed platforms, and address file system issues
- Transition GPGPU RAID work from R&D to deployed prototype for SNL use (equipment already in place), and work with LSI on next-generation RAID controller

- Continue activities for pNFS, POSIX IO API extensions, HEC/IWG FSIO
- Evaluate Hadoop/Visualization Toolkit (VTK) for *in-situ* analysis of S3D combustion code, and continue to investigate data-ingestion challenges for data-warehouse appliances
- Design and develop I/O services for shock physics, for fusion, climate, and combustion applications
- Explore modifications to 9p networking protocols to support HPC I/O
- Develop SST simulators for disk and SSDs
- Begin work on NVRAM simulators

Expected deliverables in FY11:

- Prototype in-transit analysis services for CTH shock physics code and XGC fusion code
- Production-level release and deployment of GPGPU RAID
- Prototype disk and SSD simulators for SST
- Preliminary evaluation of Ceph testbed for SNL computer operations
- Report on viability of Hadoop/VTK for in-situ analysis of ASC codes

Preliminary planned activities in FY12:

- Continue to address I/O scalability and reliability issues on Cielo
- Work with LANL to address I/O issues for the next ASC capability system
- Explore system-software to support I/O services in a production environment

WBS 1.5.4.6: Post-Processing Environments

This level 4 product provides integrated post-processing environments to support end-user visualization, data analysis, and data management. The scope of this product includes planning, research, development, integration and deployment, continuing customer/product support, and quality and reliability activities, as well as industrial and academic collaborations. Projects and technologies include tools for metadata and scientific data management, as well as general-purpose and application-specific visualization, analysis, and comparison. Research includes innovative data access methods and visualization of massive, complex data—the use of open-source foundations will continue to be an important strategy for development of shareable advanced techniques. The product must develop solutions to address interactivity, scaling, tri-lab access for petascale platforms, and data analysis techniques needed to support effective V&V and comparative analysis. Solutions for emerging platform architectures may in turn require customization and/or re-architecting of software to leverage hardware features. A continuing emphasis will be placed on tools for improving end-user productivity. The product also provides and supports infrastructure including office and collaborative space visualization displays, mechanisms for image data delivery, and graphics rendering hardware.

Post-Processing Environments Deliverables for FY11

- Hopper and chopper software updates deployed on LLNL platforms
- Data analysis and visualization software suite deployed on LLNL clusters

- New weapons science visualizations
- An operational post-upgrade Viewmaster visualization cluster
- Scientific visualization support for users
- EnSight upgrades
- PoP and GMV upgrades
- A portable open-source supercomputing rendering library based on OpenCL
- Exploration of possible integration of statistical sampling and visualization/analysis to provide a mathematically rigorous approach to effectively reducing massive data sizes
- Development of a lightweight analysis interface for line-outs to support code verification
- Develop ParaView and Ensign readers, as well as additional visualization and analysis functionality
- Remote analysis of simulation results and ensembles of runs through tools built on open source foundational technologies (ParaView for scientific analysis and browser-based analysis of ensembles)
- Ensemble analysis tools incorporating uncertainty analysis capabilities

WBS 1.5.4.6 Scientific Visualization (LLNL)

The Scientific Visualization project develops and supports tools for managing, visualizing, analyzing, and presenting scientific data, as well as conducts research. Research includes topological analysis, particle visualization, and data compression techniques. Operational support for data analysis covers support of post-processing resources, including visualization servers, displays, software, and facilities. The visualization hardware architecture team engages in planning, test bed prototyping, testing of systems and components, and procurement and integration of new systems. The operational team's server efforts include system administration, computer security, troubleshooting, and maintenance of hardware and software. Display efforts include support of high-resolution, high-performance display devices for theaters and collaborative use areas. Operational support manages theater and PowerWall facilities and associated servers, runs video production labs, and consults on software such as resource management tools, movie players, animation, and visualization packages.

The project exploits the latest capabilities of clustering hardware, graphics processing unit advances, and parallel storage systems. Hardware capabilities include three production visualization servers and several PowerWall clusters. A video display infrastructure drives power walls and smaller displays. The project installs, maintains, and consults on software visualization tools, and supports demonstrations on the PowerWall. The project maintains unclassified and classified video production labs with video editing, 3D modeling and animation tools, and video peripherals for DVDs and videotapes.

In FY10, the LLNL visualization team explored new architectures for driving PowerWalls with smaller numbers of more powerful nodes. The team also evaluated technologies for remote display, procuring and testing PC-over-IP (PcoIP) technology. The team delivered a suite of visualization software on the new SCF cluster, Graph. The team supported usage of the theaters, visualization, and video laboratories, as well as

maintained the tool suite and provided troubleshooting for software issues. Visualization researchers continued to perform work in areas of topology, compression, and advanced data analysis techniques.

Planned activities in FY11:

- Perform R&D in topological analysis, data compression, and particle visualization, and continue to mentor students in these areas
- Support petascale data analysis hardware upgrade on the OCF
- Support petascale data analysis activities through software consulting and research activities
- Maintain and enhance the existing suite of tools and libraries developed by this project in support of PowerWall and visualization cluster usage
- Provide operational support for all visualization facilities, including supporting projection equipment, performing color alignment and projector alignment, and facilitating the use of the data analysis clusters and associated storage
- Support ASC scientists through visualization and video efforts, including the support of PowerWall presentations, creation of visuals and movies to support the presentation of scientific data, and general user support

Expected deliverables in FY11:

- Enhanced remote rendering options from visualization servers
- Suite of analysis and visualization software deployed on new OCF IDA cluster

Preliminary planned activities in FY12:

- Continue to enhance and maintain the data analysis environment in both hardware and software areas, with a focus on exascale challenges
- Leverage research efforts in visualization and data analysis and perform new research in data compression and topological analysis for scientific computing

WBS 1.5.4.6 Scientific Data Management (LLNL)

In the age of petascale computing environments, the complexity and scale of the data management challenge is also reaching new heights. The Scientific Data Management project provides users with powerful and time-conserving ways to access, search, compare, and archive large-scale scientific data. This is achieved through the development of production-quality applications that enhance existing data management tools as well as provide new and innovative capabilities.

The Scientific Data Management tools team has expertise in data transport protocols, graphical user interfaces, Web technologies, data representation, databases, and advanced system architectures. The team has decades of combined experience in designing and developing productivity-enhancing applications.

Hopper and Chopper are the principal products of this effort. Hopper is a Java-based file management tool that allows users to transfer and manipulate files and directories by means of a graphical user interface. Users can connect to and manage local and remote resources using all file transfer protocols supported by ASC computing centers. Chopper is the command line version of the tool, useful in particular for automated, “background” file manipulations initiated from within applications.

In FY10, LLNL released versions 2.4 and 2.5 of Hopper and Chopper, featuring a graphical disk usage view, an aggregated copy feature for efficiently transferring vast amounts of data, a toolbar editor, and various enhancements for simplifying user interactions with production computing platforms.

Planned activities in FY11:

- Maintain and enhance the existing suite of tools and libraries developed by the scientific data management project
- Investigate additional levels of concurrency, for example, for searches and recursive directory operations

Expected deliverables in FY11:

- Version 2.6 of Hopper and Chopper, featuring job resubmission capabilities, additional parallelization of large operations, and support for additional data types in the instant-preview pane
- Version 2.7 of Hopper and Chopper, featuring scalability enhancements, usability improvements, a means for building and automatically updating site configuration files, and a history editor

Preliminary planned activities in FY12:

- Continue to maintain and enhance the Scientific Data Management suite of tools and libraries based on user feedback and requirements imposed by a petascale computing environment
- Investigate the sharing of code components with an HPC dashboard

WBS 1.5.4.6 Visualization and Insight for Petascale Simulations Project (LANL)

The Visualization and Insight for Petascale Simulations Project develops new visualization algorithms and systems to meet capability requirements for ASC petascale simulations. This work is required to address ASC workloads—massive data sizes, ensembles of results, and using unique supercomputing architectures.

The project focuses on petascale interactive visualization and analysis, which encompasses identifying appropriate hardware resources to support petascale visualization, working on data reduction based software techniques, such as intelligent data streaming, to reduce the volume of data need to be moved off the petascale platform, and exploring the use of visualization kernels that run on the petascale platform (including data analysis, visualization, and rendering methods). Running on the petascale platform is important for two reasons: 1) to get around the latency limitations imposed by the architecture (for example, if one can analyze as the data is computed, one can avoid writing to disk and then re-reading the data); and 2) for performance, using the power of the petascale platform to analyze petascale data.

In FY10, LANL made significant progress on the “visualization on the platform” milestone to support the Joint Exascale Initiative (JEI), by delivering evaluation results documented as a report and as tools made available to the ASC community. Specifically, high-performance rendering software is now available in the latest release of VTK and ParaView. Research on new petascale visual analysis approaches was published, including work on statistical, feature, and analysis approaches. New readers for Ensign and ParaView were developed for applications of interest to the ASC Program, including VPIC and POP reader support.

Planned activities in FY11:

- Develop visualization, analysis and rendering software aligned with emerging supercomputing platforms to advance petascale data analysis infrastructure
- Explore and evaluate the maturity level of visualization and rendering on the supercomputing platform

Expected deliverables in FY11:

- Develop a portable open-source supercomputing rendering library based on OpenCL
- Research the integration of statistical sampling and visualization/analysis to provide a mathematically rigorous approach to effectively reduce massive data sizes
- Develop a lightweight analysis interface for line-outs to support code verification
- Develop ParaView and Ensign readers and additional visualization and analysis functionality for applications of interest to ASC program, such as the Data Restructuring in a Production Application milestone

Planned activities in FY12:

- Transfer technology to the production environment, based on a positive evaluation of the FY10 Level 2 milestone *Visualization-on-Platform Technology*
- Explore the end-to-end integration of *in-situ* analysis, feature-extraction, and ray-tracing based visualization to improve the understanding of massive petascale results

WBS 1.5.4.6 Production Systems for Visualization and Insight Project (LANL)

The primary goal of this project is to provide LANL weapons designers with visualization systems research and support and fully utilize the systems provided. Visualization and visual analysis are essential tools needed by code teams and designers in understanding the terabytes of data that are generated in a single simulation run. This project, also referred to as “Production Visualization,” provides visualization services from the machine to the desktop for users in the ASC Program.

Technical staff members funded by Production Visualization assist in the design and deployment of new machines from participating in design teams, to overseeing the procurement of such systems, to troubleshooting graphics systems on these machines, to performing the visualization software integration tasks needed, and by installing and maintaining critical visualization software on ASC machines. They also support and maintain LANL’s large visualization facilities, including the CAVE, the Powerwall Theater, and the co-laboratories. They maintain the visualization infrastructure, which delivers video from the machines to the users’ desktops.

The project also has a small effort to maintain and support PoP and GMV, and to port these custom visualization and analysis codes to the new ASC platforms. The two staff most experienced with these two codes and their associated link files work part time on this effort while also collaborating with the ASC code teams in the weapons program.

The project’s capabilities include the design and deployment of new visualization systems, briefing support, and support of large facilities, such as the CAVE and the PWT. The project provides assistance, training, and developing of new tools to work

with these facilities. Development, deployment and maintenance of any needed Visualization Corridor software is also provided by this project.

In FY10, LANL added advanced python scripts, which now enable comparison of turbulence simulations to standard turbulence models.

Planned activities in FY11:

- Upgrade Viewmaster visualization cluster
- Participate in Cielo operational activities; in particular, the integration of visualization capability for all three labs into the Cielo platform
- Support users in providing visualizations of simulations, experimental data, and engineering computer aided design (CAD) data on desktops and in facilities
- Continue support and maintenance of the large visualization facilities
- Provide facilities support for briefings, as needed
- Analyze networking requirements for remote visualization, in cooperation with a networking requirements study
- Provide contract management and requirements specification, including facilities, visualization cluster, and EnSight contracts
- Direct EnSight development activities under the new LANL EnSight development contract to Computational Engineering International

Expected deliverables in FY11:

- An operational post-upgrade Viewmaster visualization cluster
- Scientific visualization support for users
- High availability in the large visualization facilities
- EnSight, PoP, and GMV upgrades

Preliminary planned activities in FY12:

- Upgrade visualization facilities
- Continue to support and maintain production visualization systems

WBS 1.5.4.6 Physics-Based Simulation Analysis Project (LANL)

The purpose of this project is to help LANL weapons designers utilize the full power of the hardware and software infrastructure for visualization and data analysis developed and deployed by ASC, to improve the physics understanding of their weapons simulations. To achieve this goal, the project has deployed within the design community in X Division, a small group of individuals with expert knowledge in both visualization and weapons science to work directly with the designers. The job of this small group of experts is to help designers apply the full potential of the ASC visualization and analysis infrastructure to solve their analysis problems and to promote new weapons science discoveries using the ASC codes.

This small group of individuals has played a major role in successes of the Thermonuclear Burn Initiative (TBI) program at LANL. Their work in combining ASC visualization with the new capabilities provided by the ASC codes has been one of the major factors responsible for the new discoveries being made by the TBI program. For example, the project has demonstrated for the first-time a detailed comparison between

LANL and LLNL ASC code results from a 1.76 billion cell 3D TBI simulation, which was performed on the ASC Purple machine. The new 3D visualizations and analysis capabilities for TBI simulations were used as an important part of the briefing given to Secretary of Defense Robert Gates in his recent visit to Los Alamos.

In addition to working directly with the design community on its visualization and analysis problems, this group, along with people from the Production Systems for Visualization and Insight project, is responsible for some LANL activities related to the EnSight visualization and data analysis software. This includes maintaining the EnSight software installation lab-wide, providing local user support in the use of the software and acting as a bridge between the LANL design community and the EnSight developers at Computational Engineering International for problem reporting and resolution and for new feature requests. This group, with the Production Systems for Visualization and Insight group, also directs all subcontracts that LANL has with Computational Engineering International related to new EnSight development and to on-site training and consulting.

In FY10, parallel hardware volume rendering was developed by CEI. It is now a feature of EnSight available to all LANL users as part of the package.

Planned activities in FY11:

- Work directly with designers in physics-based, iterative discovery process using Petascale Visualization and Data Analysis enabled tool (EnSight)
- Support and maintain the EnSight software and help direct EnSight development activities under the new LANL EnSight development contract to Computational Engineering International
- Document our work with joint publications co-authored with X Division designers on a variety of weapons science topics

Expected deliverables in FY11:

- New weapons science visualizations
- Other visualizations and data analysis products to address a variety of DSW-related activities including the closing of outstanding SFIs

Preliminary planned activities in FY12:

- Continue to promote new discoveries in weapons science by advanced applications of visualization and data analysis in programs such as TBI
- Continue to document the results with classified papers and publications on weapons science topics jointly co-authored with X Division designers

WBS 1.5.4.6 Scalable Data Analysis (SNL)

This project delivers advanced customer-centered production analysis capabilities within an open source framework, delivered within a variety of tools targeted for specific problem domains. Foundational capabilities are released in the open source VTK, and are delivered in a range of scalable tools appropriate to the specific need being addressed. Our scalable tools allow investigation of data on a variety of platforms—everything from a laptop to a cluster. This allows users to interact with their data, whether it fits on a PC or is located on a remote cluster in another state. This vision extends to data analysis for Exascale platforms, which will require a rich set of flexible tools that can be co-designed for specific Exascale problems.

We leverage SNL's extensive work in complex data analysis in non-scientific regimes to promote analysis of large collections of data. In particular, understanding high dimensional, feature-rich data requires highly interactive and iterative investigation of the data. In addition to providing these advanced capabilities, this project provides deployment and support services that enable ASC customers to carry out data analysis on Cielo and the next ACES platform.

In FY10, SNL advanced hardware solutions for large data analysis in support of ACES analysis goals, including support for efficient rendering on the computation platform (Cielo) and mitigation of I/O bandwidth bottleneck through *in-situ* analysis capabilities. SNL released the initial implementation of tools for analysis of ensembles of runs, promoting interactive investigation of entire sets of runs as a single dataset. Current tools focus on sensitivity analysis of input parameters and output results, and are approved for use on classified data. The effort on *in-situ* analysis capability was expanded to include release of an *in-situ* library and outreach to the simulation community. This capability provides a general analysis toolset (including the fragment and crater characterization work from milestones in FY08, FY09, and advanced rendering from FY10) that can be used in conjunction with a running simulation to provide a rich set of data output on extreme size systems.

Planned activities in FY11:

- Provide targeted analysis capabilities to promote an integrated user experience to improve productivity and insight into results, in conjunction with SNL codes
- Continue research into analysis of complex data ensembles, with an emphasis on visualizing and understanding uncertainty, both in inputs (for example, constituent models) and the resulting ensembles of output results
- Partner with SNL customers to provide advanced analysis techniques through easily accessible browser-enabled interfaces
- Provide ongoing technical and user support, including direct analysis support for capability class systems

Expected deliverables in FY11:

- Remote analysis of simulation results and ensembles of runs through tools built on open source foundational technologies, for example, ParaView for scientific analysis and initial browser-based analysis of ensembles
- Ensemble analysis tools incorporating uncertainty analysis capabilities
- Advanced analysis and visualization library for direct coupling with codes
- Ongoing releases and installations of ParaView

Preliminary planned activities in FY12:

- Deliver advanced analysis capability in support of GA of the Cielo platform, including high-performance distance visualization, platform-specific optimization, and initial *in-situ* analysis capability

WBS 1.5.4.7: Common Computing Environment

The goal of the CCE product is to enable a common environment across the tri-labs that will initially be deployed on the TLCC systems. The scope of this product includes

funded R&D projects to address gap areas identified by the tri-lab technical working groups.

The CCE working groups and projects focus on a common software stack, including but not be limited to, operating system software; application development tools; resource management; HPC monitoring and metrics; and common tri-lab environment issues such as configuration management, licenses, WAN access, and multi-realm security.

Common Computing Environment Deliverables for FY11

- Tri-lab Level 2 milestone “Deploy CCE capabilities for capacity computing environment”
- TOSS 2 (based on RHEL 6) and TLCC2 integration support
- OI SS deployed on Sequoia ID (Dawn) and Cielo and the Initial Component Framework for scalable performance analysis beta made available
- Expanded Workload Characterization Tool (WC Tool) with improved custom reporting interface and additional reporting capabilities; additional configuration and administration tools
- Complete Service Level Agreement (SLA) document based on services defined by the service catalog
- Increased scalability and platform expansion of STAT lightweight debugger and MPI messaging and correctness tools

WBS 1.5.4.7 Tripod Operating System Software (LLNL, LANL, SNL)

TOSS is the tri-lab software stack to run across all newly procured Linux capacity clusters, initiating with TLCC platforms delivered in FY08. The goal of the TOSS project is to increase efficiencies in the ASC tri-lab community with respect to both the utility and cost of the CCE.

This project delivers a fully functional cluster operating system (kernel, Linux distribution, IB stack and related libraries, and resource manager) capable of running MPI jobs at scale on TLCC hardware. The system is to meet CCE requirements for providing a common software environment on TLCC hardware across the tri-lab complex, now and into the future.

TOSS provides a complete product with full lifecycle support. Well-defined processes for release management, packaging, QA testing, configuration management, and bug tracking are used to ensure a production-quality software environment can be deployed across the tri-lab in a consistent and manageable fashion.

In FY10, the following new capabilities were added: released a minor update to TOSS (version 1.2-3), which included security updates and bug fixes; developed/ deployed a TOSS major update (version 1.3); and developed a beta release of TOSS 2.0, which is based on the next major RedHat release, RHEL 6.

Planned activities in FY11:

- Provide ongoing TOSS software development and support
- Develop/ deploy GA release of TOSS 2.0, contingent on RedHat RHEL 6 schedule
- Prepare for deployment of the next generation of the ASC TLCC systems (TLCC2), which may include software integration and testing for the tri-lab environment

- Evaluate/develop new technologies

Expected deliverables in FY11:

- TOSS 2.0 deployed
- TLCC2 systems integration support provided by TOSS

Preliminary planned activities in FY12:

- Provide ongoing TOSS software development and support
- Develop/deploy TOSS 2.X (based on RHEL 6.X)

WBS 1.5.4.7 Open | SpeedShop (LLNL, LANL, SNL)

O|SS provides a wide range of performance experiments within a single environment. This includes support for PC sampling, inclusive and exclusive execution times for routines, hardware counters, as well as I/O and MPI tracing. O|SS provides all functionality through three fully interoperable interfaces that provide the greatest possible flexibility to users and code teams. Further, O|SS is designed to work on binaries of application without the need to recompile, enabling a clean and easy integration into the development workflow.

In FY10, the project continued deployment and support for ASC users and code teams. In addition, FY10 marked the start of a joint NNSA/DOE Office of Science (ASCR) project on O|SS for which this CCE project provides the necessary funding. As part of this project, we started rewriting the O|SS framework for increased modularity and scalability as well as porting O|SS to existing NNSA capability machines, with the goal of ultimately providing the toolset on future NNSA systems, including Sequoia and Cielo.

Planned activities in FY11:

- Provide production version of O|SS on the Sequoia ID (Dawn) and Cielo
- Integrate binary instrumentation
- Transform O|SS framework towards more modularity and scalability
- Improve scalability of data collection and analysis
- Provide new data collection infrastructure enabling the use of preload instrumentation with online data gathering and aggregation

Expected deliverables in FY11:

- Toolset running on Sequoia ID (Dawn) and Cielo
- Back port of all new additions for capability systems to TLCC clusters
- New hardware counter experiments following users' demands
- Initial Component Framework for scalable performance analysis constructed from existing O|SS components

Preliminary planned activities in FY12:

- Continue collaboration with OASCR on O|SS developments
- Complete implementation of new, scalable tool framework
- Provide additional analysis capabilities exploiting online aggregation

WBS 1.5.4.7 Workload Characterization (LLNL, LANL, SNL)

The WC Tool provides the capability to collect and report both current and future requirements for compute resources with programmatic characterization of the work, and also computing resource usage with programmatic characterization of the work. The tool includes a modular mechanism that can interface to multiple existing databases at each of the tri-labs. Additional development is needed in the areas of validation, test suites, and common reporting capability.

Resource Management/Moab development is used to tie WC Tool demand/estimates, with their respective workload characterization, to job requests and resulting platform usage data.

In FY10, the project improved the tri-lab WC Tool's estimate interface and its ability to collect, analyze, and report compute resource requirements with programmatic characterization of the work. During FY10, the initial reports for platform usage continued to be available for tri-lab quarterly reports. A new, custom reporting interface was developed to generate additional computing resource usage and machine utilization reports. This work included integration with the underlying platform usage databases (for example, the SLURM accounting database).

Planned activities in FY11:

- Enhance the WC Tool to improve the custom reporting interface and provide additional tools for configuration and administration
- Gather requirements from NNSA/HQ for each laboratory to expand the reporting capabilities
- Continue investigation in shared capabilities for data collection, storage, and reporting
- Support the tri-lab CCE Level 2 milestone

Expected deliverables in FY11:

- Expanded WC Tool with improved custom reporting interface and additional reporting capabilities, additional configuration, and administration tools

Preliminary planned activities in FY12:

- Assess usage and provide enhancements in tools or documentation and training
- Improve the WC Tool in coordination with the other laboratories in the context of the CCE projects

WBS 1.5.4.7 Distributed Data Services (LLNL, LANL, SNL)

To date, application, system, and platform information have been collected with varying degrees of fidelity. The information gathered has largely been stored in log files that are only viewed upon suspicion of error and then either manually or by scripts, tailored to look for particular indicators. This type of post run analysis and troubleshooting is very inefficient and will not scale with ASC systems. This Distributed Data Services (DDS) project seeks to define a common scheme for data organization and access such that the aforementioned types of information for ASC systems will not only be stored but will be accessible in a standardized and dynamic way at the fidelity desired and on time scales commensurate with collection. Such a system will allow cross site, cross platform development of mechanisms to take advantage of this new capability in a dynamic and

automated fashion, thus increasing the productivity of the platforms, users, and administrators.

This project is new for FY11.

Planned activities in FY11:

- Implement and deploy services on at least one TLCC system per lab for testing and hardening
- Develop extended Data Catalogue specifications and standards
- Investigate use of Inter-Site HPC (IHPC) environment for cross lab use of services
- Support at least one intelligent data analysis project to take advantage of the utility of access to multiple data sources through DDS

Expected deliverables in FY11:

- A prototype application deployed on at least one TLCC system at each of the three labs and using DDS services related to that system
- DDS approved for classified operation at each lab

Preliminary planned activities in FY12:

- Work with developers, code maintainers, and system administrators to augment or adapt subsystems to better interoperate with SWISS
- Measure effectiveness of FY11 DDS deployments
- Track multi-agency efforts such as CIFTS (which targets low volume hierarchical information sharing) for potential exascale enabling extensions

WBS 1.5.4.7 Shared Work Space (LLNL, LANL, SNL)

The Shared Work Space project is the infrastructure for promoting collaboration across the laboratories. It currently includes the Gforge server that is housed and managed at SNL. Gforge administrative support for the environment is done through identified staff at each lab. Gforge user licenses are purchased through LANL. During FY10 collaborative usage of Gforge by CCE Working Groups and projects increased significantly, for document management, maintaining Wikis, managing SVN repositories, and tracking project activities.

Expectations for FY11 are to continue support for the Gforge server and licenses, and initiate additional services to support tri-lab account management to set the stage for additional cross-lab services.

In FY10, Shared Workspace addressed the initial difficulties in providing cross-lab access. Users from all labs now have Gforge access via cross-realm authentication. CCE Working Groups and project teams are using Gforge for document management and other activities. The previous Tripod Wiki site is being phased out in FY10, with most content being archived in Gforge.

Planned activities in FY11:

- Maintain Gforge server for CCE project collaboration
- Define services (resources) to be offered
- Develop SLAs

- Develop requirements with input from all three labs
- Evaluate potential tools and decide on implementation methodology

Expected deliverables in FY11:

- Continue support for the Gforge server and licenses
- Shared Workspace Service Catalog (including SLAs)

Preliminary planned activities in FY12:

- Continue maintenance and increased use of Gforge and other Shared Workspace services

WBS 1.5.4.7 Gazebo Test and Analysis Suite (LANL, SNL)

Gazebo is a suite of software components used to test, monitor, and analyze the health of an HPC system. The Gazebo tool also includes the Resource Health Characterization Engines (RHCE), a set of statistical analysis algorithms that model and identify anomalies within the characteristics of metrics associated with an HPC platforms.

The Gazebo software suite provides a framework where system test data is created and stored in flat files and on a networked accessible database. Enhancements to the data included within each test, in the form of general and specific performance and correctness data, is saved for later analysis. This data by itself can and is used to judge the health of an HPC system given its trending over a period of time. Currently, data is being generated and collected. Also, a Web interface based on the Model/View/Control paradigm written in Ruby/Rails provides a visual display of test results.

In FY10, the Gazebo project integrated the RHCE component. This software provided a smart mechanism to normalize the results of performance based tests and learned what is a viable set of benchmarks. Users can now run a test and have the system learn what is normal, leading to detection of anomalous system behavior. Also, updates to some of the system interfaces were changed per requests from users at participating labs.

Planned activities for FY11:

- Provide input (for example, database schema design prototype, user interface requirements, and data protection requirements) to support the DDS effort.
- Participate in design goals and tool discovery efforts of the DDS project with a Gazebo focus (this development and integration thrust is a significant step towards a goal of complete system monitoring and diagnosing the health of HPC systems)
- Perform system testing

Expected deliverables in FY11:

- A working proof-of-concept data-sharing model that allows tool developers within the monitoring working group the ability to share their data using a common interface (each stage of the deployment, including architecture design, tool discovery, and showing a basic working model)
- Specific design criteria necessary to incorporate Gazebo data into the DDS cloud
- Identification of and provided specific data for use in a DDS prototype
- Support for initial testing of TLCC2 procurement expected in FY11–FY12

Preliminary activities for FY12:

- Provide full integration of all the tools within the CCE space

WBS 1.5.4.7 Debugger (LLNL, LANL, SNL)

All three labs are utilizing TotalView as the core debugger. LLNL has the more strategic relationship with TotalView Technologies and is also doing additional research into other debugging tools (for example, STAT). The approach to subset debugging is key as ASC moves toward larger scales and collects information that helps focus in on the trouble area. Input from all three labs and target applications are required. The working group would like to build a tri-lab debugger capability around the LLNL capability.

The focus of this project is to develop a more cohesive debug capability that is in line with the CCE goals of commonality and leverage of infrastructure and plans. This includes a range of sub-projects that include:

- Further develop the communication/support subteam led through LLNL.
- Create a common repository of debugger tool documentation hosted on the CCE Gforge server.
- Provide a central integration point for TotalView strategic planning through LLNL.
- Provide cross-lab deployment support for STAT. This is focused on subset debugging that would also utilize TotalView and support scaling efforts. The University of New Mexico will provide technical guidance on porting and adapting the base infrastructure.
- Increased scalability and better integration of MRNet, which is the basis for STAT, through collaboration with the University of New Mexico.
- Work towards addressing debugging on capability machines through the subset-debugging segment of this capability.

In FY10, the project was able to deploy STAT on TLCC machines at each of the tri-labs and both simplify the build process and initiate efforts to port to non-SLURM machines. In addition, both sharing of information and plans regarding TotalView deployment were initiated as well as an initial cross-lab repository for documentation and training material.

Planned activities in FY11:

- Further hardened and expanded for larger scale via STAT (data transport)
- Port launching for non-SLURM machines and develop enhancements for usability based on user needs via STAT
- Continue coordination for issues, strategic needs via TotalView
- Develop the process and infrastructure to share documentation, training, issue resolution, and support

Expected deliverables in FY11:

- Increased scalability and platform expansion of STAT components (MRNet and integrated to additional launching mechanisms)
- Increased debugger tri-lab support approach
- Further deployed TotalView and STAT

Preliminary planned activities in FY12:

- Increase the capability to scale debugging tools
- Increase the capability to provide users debugging tools that increase productivity

WBS 1.5.4.7 Inter-Site High Performance Computing Security Integration (LLNL, LANL, SNL)

With the advent of the TLCC effort, LANL, LLNL, and SNL have an increased need for integrated access to unclassified resources for improved collaboration and to better utilize computing resources. Attempts to facilitate such collaborations have been hampered by incompatible security policies, interpretations, and implementations that make it difficult to access unclassified resources cross-site. These issues spawned the inter-site HPC integration effort. The project has developed an architecture based upon a Tri-Labs Interconnection Security Agreement (ISA) and a governance model. Tri-lab inter-site HPC service deployment and network enhancements are the target for this year.

The target capability is to identify and implement a tri-lab security approach to achieve the following targets:

- Security policies and implementations that allow tri-lab access using identified protocols and technologies
- More effective resource control and utilization
- Utilization, where feasible, of home site security apparatus, including authentication equipment (for example, single sign on)
- User access to tri-lab resources within a specified period of time upon request based on identified need
- Collaboration with the CCE user community to develop security policies that facilitate easier resource utilization
- Collaboration with the appropriate security personnel to ensure timely implementation
- Support of multiple classes of users, such as architecture and modeling, application, customer service, system administrators with root access, and end users
- Necessary network bandwidth and latency to cover current and projected requirements

In FY10, the inter-site HPC architecture was developed and the tri-lab governance model was agreed upon. This required much discussion and compromise to consider the differences at the various sites and local security policies. Project planning started in FY10 and continues in FY11.

Planned activities in FY11:

- Complete installation of 1 Gb networking infrastructure based on ISA design
- Complete SLA document based on services defined by the service catalog
- Deploy Services defined by Service Catalog
- Upgrade networking infrastructure to 10 Gb

- Support the tri-lab Level 2 milestone “Deploy CCE capabilities for capacity computing environment”

Expected deliverables in FY11

- DAA approval of ISA
- Tri-lab approval of Service Catalog

Preliminary planned activities in FY12:

- Integrate collaborative and other services between sites

WBS 1.5.4.7 Open Source Contract Maintenance (LLNL, LANL, SNL)

This project provides funding to outside developers who maintain tools and tool infrastructures that are critical for code teams or serve as the basis for internal tools. This funding will be provided to those developers through support contracts administered by LLNL, but each contract includes support for all three laboratories and all three laboratories in close collaboration provide the technical guidance for the three contracts.

In FY10, the project provided a structured process to provide support for some of the critical open source tools and tool components. During FY10, this project provided funding for OI SS, Valgrind, TAU, as well as supported workshops at three labs to increase visibility of ASC/CCE tools in the tri-lab user community.

Planned activities in FY11:

This project aims at providing the following activities for each of the three projects:

OI SS, Krell Institute

- Provide support in the form of support calls and code team interactions
- Provide training and tutorials at all three laboratories (one per lab)
- Expedite bug resolution and patch releases for the labs
- Perform continuous testing and release validation on lab systems

Valgrind, OpenWorks

- Provide support for tri-lab platforms (capacity and capability)
- Perform release validation on lab systems
- Expedite bug resolution and patch releases for the labs

TAU, Paratools

- Provide support in the form of support calls and code team interactions
- Provide training and tutorials at all three labs
- Coordinate among the tri-labs for contract execution

Expected deliverables in FY11:

- Support and maintenance for the three open source software components
- Quantifiable number of support hours delivered to the tri-labs
- Tutorials and/or workshops held at all three labs

Preliminary planned activities in FY12:

- Reevaluate which software components require support contract
- Continue support and maintenance for open source software components
- Coordinate among the tri-labs for contract execution

WBS 1.5.4.7 OpenMPI Integration/Scaling (LLNL, LANL, SNL)

With the increasing complexity of planned capacity and capability systems and the growth in scale, it becomes important to establish a tri-lab effort that can assess MPI functionality and performance to enact changes supporting the ASC workload. OpenMPI is used as the primary MPI at LANL and SNL and as the secondary MPI at LLNL. MVAPICH is the primary MPI at LLNL with OpenMPI as the secondary. These factors mandate a continued need for better integration with the resource manager, for scaling studies, and investigations for future resilience improvements. LANL is a member of the OpenMPI development community and can effect improvements targeted at capability systems. The addition of tri-lab testing, performance monitoring, and support will allow better application efficiency through parameter studies on ASC clusters against various application usage models.

This project is targeted to develop a set of capabilities focusing on supporting scale increases, assessing performance of both MPI and user applications, and providing optimal parameters to users for better MPI performance. The complex multi-socket, multi-core NUMA node architecture of systems like TLCC mandates such an investigation. Close working relationship with the OpenMPI community and other MPI developers is seen as a strategic need.

In FY10, the project provided both optimizations to OpenMPI and cross MPI support for tools and scalability. The following capabilities were added: OpenMPI; integration of wire-up improvements into the production version of OpenMPI; bug fixes to OpenMPI for tri-lab issues, Panasas checking, and other issues; system V shared memory component for intra-node communication; MPI Scalability Support; prototype scaling test harness tested on tri-lab machines; assess MPI Correctness tools efforts such as MUST; and beta release of Loba tool to collect user application communication patterns.

Planned activities in FY11:

- Broaden availability and improve deployments of OpenMPI and MVAPICH
- Assess MPI performance, identify best build and run parameters through studies, investigate impact of topology mapping, and provide results to end users
- Collaborate with MPI Correctness tools efforts such as MUST
- Provide enhancements/bug fixes to OpenMPI/ROMIO decoupling

Expected deliverables in FY11:

- Performance studies across different MPIs, OpenMPI/MVAPICH/ other MPI implementations
- Tools to support MPI at increased scale
- Modifications and enhancements to OpenMPI based on tri-lab need
- User tools to customize application run parameters

Preliminary planned activities in FY12:

- Continue to build tri-lab capability, scaling assessment tools, performance enhancements geared toward tri-lab applications

WBS 1.5.4.7 File System Architecture and Integration (LLNL, LANL, SNL)

Both LANL and LLNL have developed I/O middleware to speed up I/O for codes that run on their supercomputers. LANL developed the PLFS and LLNL developed the Scalable/Check-Point Restart (SCR) code.

A substantial amount of time and effort has been put into each of these I/O mechanisms and both have shown benefits that can be realized in a production setting. This project will enable for both PLFS and SCR to be further introduced to the tri-lab HPC environment. This would include tri-lab efforts to install on specific machines for assessment, modularize both code for a common standard API, such as POSIX, to require no application changes to each, work out issues involved in a wider production installation and support process, and form the basis for working on a file systems architecture on which these and other projects can build.

PLFS is a virtual parallel file system that reorganizes logical access to a single parallel file into physical access to multiple, non-shared files (for example, it converts N-1 parallel IO into N-N). This has been demonstrated to improve checkpoint bandwidth by at least an order of magnitude on the three major parallel file systems: Lustre, GPFS, and PanFS.

SCR is an effort at LLNL with an identical motivation—to improve checkpoint bandwidth by eliminating file system contention. Additionally, SCR makes use of local storage devices by saving checkpoint data directly on the compute nodes' disk drives or to the nodes' memory.

This is a new project for FY11.

Planned activities in FY11:

- Implement PLFS/SCR on machines at the three labs for assessment of installation process and capability to support local environment and applications; iterate findings with development lab
- Support integration into TOSS
- Hold tri-lab discussions toward file system architecture approaches that support lab needs

Expected deliverables in FY11:

- Tri-lab installments of PLFS/SCR
- Architecture document in file systems area

Preliminary planned activities in FY12:

- Implement PLFS/SCR as found appropriate in FY10
- Continue work on tri-lab file systems collaboration

WBS 1.5.4.7 SARAPE Re-Development (LLNL, LANL, SNL)

As part of the Shared Work Space environment, it is desirable to offer a Service Catalog through which collaborators can view and request accounts and services available in the shared environment. We propose to provide this catalog by significantly improving and expanding the Synchronized Account Request Automated ProcEss (SARAPE) tool currently used to manage resource requests within the Nuclear Weapons Complex. The current version of SARAPE processes hundreds of requests for over 50 resources at the three host laboratories, in a few days, as contrasted to several weeks or months previously required using paper forms and manual handling. It would be a natural extension of this tool to offer tri-lab collaborators additional services and resources in a Service Catalog within the Shared Work Space.

SARAPE was developed in what is now considered a legacy Web application. We propose to evaluate alternative implementation technologies to re-write SARAPE; one potential solution consistent with other CCE tools would be using Ruby on Rails. Also included will be a significantly improved user interface, further automation of administrative and management tasks, and new service offerings. A modular design will be employed that will consistently manage CCE Shared Workspace account/service access while providing modules that can be further developed at each site to improve management of local resources.

SARAPE is a Web-based application that allows users within restricted domains to request selected CCE resources to which they are permitted access. A designated Guest Processing Agent (GPA) at the requestor's site provides verification and approval for each request. SARAPE manages each user's characteristics (work location, citizenship, clearance level) and each resource's requirements (security requirements, access parameters) and tailors resource offerings to each user appropriately. A designated Host Processing Agent (HPA) at the resource site is responsible for verification and approval of requests.

Resource request submittal within SARAPE is administratively automated to a large extent, and resource availability is parameterized such that users are offered access only to resources for which they satisfy minimum requirements. Capabilities such as this eliminate confusion by providing complete information about all resources available to a requestor, and save time by ensuring requests are correctly submitted and managed.

This is a new project for FY11.

Planned activities in FY11:

- Define services (resources) to be offered
- Develop SLAs
- Develop requirements with input from all three labs
- Evaluate potential tools and decide on implementation methodology
- Begin development of new version of SARAPE to include defined services

Expected deliverables in FY11:

- Shared Workspace Service Catalog (including SLAs)
- New version of SARAPE to include defined services (initial implementation hopefully started mid-year)

Preliminary planned activities in FY12:

- Define and deploy additional Shared Work Space services

- Develop and deploy enhanced GPA /HPA modules to allow further automation of request processing and resource approval at each site, through standard APIs
- Develop and deploy enhanced user database to assist in Identity Management for Shared Workspace

WBS 1.5.5: Facility Operations and User Support

This sub-program provides both necessary physical facility and operational support for reliable production computing and storage environments as well as a suite of user services for effective use of ASC tri-lab computing resources. The scope of the facility operations includes planning, integration and deployment, continuing product support, software license and maintenance fees, procurement of operational equipment and media, quality and reliability activities, and collaborations. FOUS also covers physical space, power and other utility infrastructure, and LAN/WAN networking for local and remote access, as well as requisite system administration, cyber-security, and operations services for ongoing support and addressing system problems. Industrial and academic collaborations are an important part of this sub-program.

WBS 1.5.5.1: Facilities, Operations, and Communications (Retired)

This level 4 product was retired at the end of FY10 and replaced with two new products, which began in FY11:

- 1.5.5.4: System and Environment Administration and Operations
- 1.5.5.5: Facilities, Network, and Power

WBS 1.5.5.2: User Support Services

This level 4 product provides users with a suite of services enabling effective use of ASC tri-lab computing resources. The scope of this product includes planning, development, integration and deployment, continuing product support, and quality and reliability activities collaborations. Projects and technologies include computer center hotline and help-desk services, account management, Web-based system documentation, system status information tools, user training, trouble-ticketing systems, and application analyst support.

User Support Services Deliverables for FY11:

- Deployment of the Front Range Problem Management module
- Sequoia ID (Dawn) system training to ASC and Alliance partners
- Ongoing customer support metrics
- Reliable and responsive service to users in the ASC tri-lab computing environments
- Customer support via tri-lab service level agreements for transferring HPC support issues among labs
- Online support on classified network
- On-going user support for ASC/LANL/ACES computing users
- Establishment of tri-lab telephone support for ACES/Cielo
- Integrated ACES problem tracking system for Cielo

- Cielo system on-line documentation
- Coordinated, tiered user support for SNL's ASC resources
- Support SNL ASC-related user support resources
- Collaborate to provide tri-lab user support capabilities and to support the Predictive Science Academic Alliance Program (PSAAP)

WBS 1.5.5.2 Hotlines and System Support (LLNL)

The Hotlines and System Support project provides users with a suite of services enabling effective use of ASC tri-lab computing resources. This project includes computer center hotline and help desk services, account management, Web-based system documentation, system status information tools, user training, incident management systems, and application analyst support. Services are provided to both LLNL users as well as users from external sites including LANL, SNL, and the ASC Alliance sites.

In FY10, LLNL deployed a new incident management system. LLNL also developed documentation and training materials for the Sequoia ID (Dawn) system.

Planned activities in FY11:

- Deploy and integrate Front Range Problem Management module
- Evaluate Front Range Change and Release Management modules
- Provide ongoing support services for hotline operations, documentation, and training
- Evaluate hotline metrics on a regular basis to identify potential areas that would benefit from improved documentation and training

Expected deliverables in FY11:

- Deployment of Front Range Problem Management module

Preliminary planned activities in FY12:

- Continue to implement ITIL best practices
- Implement Change and Release management Front Range modules
- Evaluate Front Range Configuration management module
- Provide ongoing support services for hotline operations, documentation, and training

WBS 1.5.5.2 Integrated Computing Network Consulting, Training, Documentation, and External Computing Support (LANL)

The Integrated Computing Network Consulting, Training, Documentation, and External Computing Support project is responsible for direct customer service for local and remote users of ASC/LANL resources, the development and delivery of documentation and training materials for ASC/LANL resources, usage statistics, and an administrative interface for ASC tri-lab and Alliance users, and other external ASC/HPC users. The primary capabilities consist of user support services, operational metrics for an HPC environment on, for example, usage and availability, Web page development to present

this information to system personnel and users, and the development of user documentation and training.

In FY10, LANL developed an initial user training course for Cielo and began working with SNL's emerging HPC support function to establish agreements and procedures as to how to jointly support an HPC resource. LANL also hired one of two additional full-time employees to ensure adequate staffing for upcoming tri-lab support demands.

Planned activities in FY11:

- Perform ongoing user support for users of ASC/LANL/ACES computing resources
- Continue to improve the quality of support for productive use of computing resources
- Provide a unified user support interface to all tri-lab users of Cielo and its associated development and test machines; help will be available via users preferred mechanisms, including telephone, email and Web interfaces
- Integrate ACES support mechanisms with existing internal support mechanisms
- Continue to explore alternative communication and support models for providing responsive support within the ASC program

Expected deliverables in FY11:

- Cielo system online documentation
- Tri-lab telephone support for ACES via a Cielo toll-free number
- Integration of local trouble tracking systems such that they automatically exchange information about Cielo issues
- A tri-lab solution for sharing existing and to-be-developed HPC documentation in a presentation-free manner suitable for adaptation to local documentation Web sites

Preliminary planned activities in FY12:

- Deliver Cielo system training to tri-lab users

WBS 1.5.5.2 Tri-Lab System Integration and Support (SNL)

The Tri-Lab System Integration and Support project manages projects relating to tri-lab production networking services and related infrastructure. SNL provides coordination, operational support, and oversight to develop and operate the ASC WAN and manages the communication link contracts. SNL leads the integration of new encryptor technology into the WAN by evaluating early engineering samples and organizing tri-lab-wide functional testing prior to deployment.

Traffic engineering and modeling systems, as well as a dedicated test laboratory based WAN development environment, are used to improve network efficiency and utilization. Monitoring and management systems are used to analyze network performance and validate vendor availability data to ensure proper credits are applied to the communication link contracts. The project oversees the Qwest communication link contract and monitors ESNET connectivity. It provides system-level analyst support for cross-site production services related to data transfer, distance computing, and access methods and services. The project coordinates production requests for tri-lab resources.

This project provides personnel for coordinating all tri-lab computing, networking, monitoring and debugging activity. Testing of the reliability of the ASC WAN is an

ongoing activity. Monitoring of services and servers present on the WAN provides immediate notification to support personnel who respond when errors are detected.

In addition, this project provides for support of the CCC utilization reporting and prioritization activities for ASC capability resources.

In FY10, SNL added monitoring of classified network connections to LANL in preparation for Cielo installation.

Planned activities in FY11:

- Operate tri-lab ASC WAN and manage the Qwest communication link contract
- Instantiate the IHPC unclassified collaborative environment with available Cielo resources at SNL
- Test and validate the performance of the ASC WAN for access to Cielo from SNL and LLNL
- Support the data transfer requirements for remote computing on Cielo at LANL

Expected deliverables in FY11:

- Operational environment for access to Cielo from SNL desktops using single sign on authentication
- Operational environment for unclassified collaborations between LLNL, LANL, and SNL.

Preliminary planned activities in FY12:

- Testing of high-speed encryption devices and technology exploration of new interconnect protocols or technologies to prepare for exascale systems

WBS 1.5.5.2 User Support (SNL)

The User Support Project provides user support and user support tools for SNL computing systems and tri-lab resources. User support activities focus on improving the productivity of the entire user community, local and remote, in utilizing the ASC HPC resources.

Tri-lab support includes assisting tri-lab customers with problems at SNL; assisting SNL customers with computing at remote locations; management of SNL computing resources via the CCC process for Purple and Cielo, and the SNL Platform Oversight Committee process; and representing SNL needs to the expedited priority run process. The support team works with a breadth of ASC applications and system environments to develop and apply expertise that enables efficient and effective use of ASC's computing resources.

This project deploys and maintains the following user support tools and functions for ASC user support efforts:

- Coordination between user support activities and leadership in adopting ITIL principles and practices
- ITIL-compliant incident, problem, and knowledge management tools (BMC ITSM Suite) and team-specific configuration of this tool
- Training facilities and equipment
- A Web portal for HPC-related information, real-time data, and documentation

In addition, this project provides the following user support capabilities in conjunction with other projects:

- A tiered user support structure (HPC Service Desk) that responds to SNL and tri-lab user requests received via phone, email, Web-based requests, and in-person visits
- The SARAPE tri-lab account provisioning Web-based tool
- Web-based, classroom, and one-one training
- Direct support in utilizing ASC resources

This project also funds the ASC User Support Team's involvement in collaborative efforts such as PSAAP and ACES.

In FY10, SNL added the following new capabilities: developed the tiered user support structure (the HPC Service Desk); helped purchase, provided HPC teams' configuration, and piloted the use of the ITIL-compliant Service Desk tools within the BMC ITSM Suite; and deployed the SNL HPC Web Portal (<https://computing.sandia.gov>) to production.

Planned activities in FY11:

- Provide coordinated, tiered user support for SNL's ASC resources
- Deploy user support function for Cielo, in partnership with LANL
- Support SNL ASC-related user support resources (for example, BMC ITSM Suite, training facilities, and HPC Web portal)
- Collaborate to provide tri-lab user support capabilities and to support PSAAP

Expected deliverables in FY11:

- Joint service desk for the Cielo platform with LANL ASC user support
- Coordinated, tiered user support for SNL's ASC resources
- Improvements to tiered user support structure and user support tools

Preliminary planned activities in FY12:

- Continue to improve the HPC service desk
- Adoption of additional ITIL principles and practices

WBS 1.5.5.3: Collaborations

This level 4 product provides programmatic support for collaboration with external agencies on specific HPC projects. This product also includes collaborations with internal or external groups that enable the program to improve its planning and execution of its mission.

Collaborations Deliverables for FY11

- Support for Supercomputing conference and SC10 ASC Research Exhibit
- Host two Predictive Science Panel meetings

WBS 1.5.5.3 Program Support (LLNL)

The Program Support project provides service to the ASC program. Program Support services include procurement and contracting, project management, and meeting support. These services are in support of both LLNL-only and tri-lab activities.

In FY10, the most significant change in the Program Support project was the initiation of TLCC2 procurement activities.

Planned activities in FY11:

- Continue management of the Sequoia contract and its associated R&D and D&E contracts with IBM for BlueGene/Q
- Manage existing tri-lab contracts and negotiate/execute any new contracts
- Support the annual Supercomputing conference, Predictive Science Panel meetings, and other meetings and workshops
- Finalize the Best Practices Workshop #4 activities and participate in Workshop #5 as needed
- Support PSAAP collaborations
- Provide support to the ASC Federal program-management office
- Participate in planning with the HPC community for a possible future exascale program in partnership with the Office of Science; this may include activities such as workshops

Expected deliverables in FY11:

- Execution of any new CCE and TLCC-related contracts

Preliminary planned activities in FY12:

- Continue FY11 activities

WBS 1.5.5.3 Scientific Collaborations (LLNL)

This project provides support for scientific collaborations with the Office of Science for multi-institution SciDAC projects, and with the Defense Threat Reduction Agency (DTRA) Nuclear Weapons Effects Division.

LLNL scientists, together with existing codes and resources at LLNL, have unique capabilities to address scientific challenges of interest to other parts of DOE and OGAs.

In FY10, LLNL provided project reports directly to the lead collaboration agencies.

Planned activities in FY11:

- Participate in a joint NNSA / Office of Science SciDAC project led by the University of Southern California to develop a hierarchical petascale simulation framework addressing stress corrosion cracking in metals and alloys from first principles, with specific emphasis on the hybrid coupling of quantum simulations and quantum-based atomistic simulations to develop optimized potentials for stress corrosion cracking applications. Detailed project planning and accomplishments are reported separately to DOE through the project principal investigator.
- Participate in a joint NNSA / Office of Science SciDAC project led by Stanford to develop improved numerical methods for flows involving shocks, turbulence, and strong density gradients with special emphasis on ensuring that the new methods

scale to hundreds of thousands of processors. Detailed project planning and accomplishments are reported separately to DOE through the project principal investigator.

- Continue to collaborate with DTRA Nuclear Weapons Effects Division to investigate electromagnetic pulse signature of several new types of electromagnetic pulse-generating phenomena and begin studies of several new candidates for observed, but unexplained particle jetting to satellite-relevant altitudes.

Expected deliverables in FY11:

- Project progress reports to the lead collaboration agencies

Preliminary planned activities in FY12:

- Continue FY11 activities

WBS 1.5.5.3 Program Support (LANL)

Through the Program Support project, LANL provides support to the national program, both by providing resources and expertise to the Federal program office and by participating in coordination and integration activities for the tri-lab program.

In FY10, LANL hosted the Predictive Science Panel, provided consultant support to the Federal program efforts to foster collaborations while building support within the predictive science community, supported the PSAAP, and provided support for the Supercomputing conference.

Planned activities in FY11:

- Alternate with Livermore in hosting the Predictive Science Panel; results will be incorporated into program plans and initiatives
- Provide consultant support to the Federal program management efforts to foster collaborations and build support within the predictive science community
- Support the PSAAP

Expected deliverables in FY11:

- Organization of the ASC tri-lab booth at the SC10 conference
- Predictive Science Panel meeting

Preliminary Planned activities in FY11:

- Host the Predictive Science Panel

WBS 1.5.5.3 Collaborations (SNL)

The Collaborations project funds critical coordination and integration activities essential to the success of ASC. These are divided into two distinct parts: 1) provide ASC multi-level communications per existent communications plan and by special request, and 2) SNL outreach to the Department of Defense (DoD) laboratories and programs.

SNL held the second meeting of the Predictive Engineering Science Panel, an external review of integrated simulation and phenomenological science activities crucial to the establishment of predictive capabilities for engineering assessment and certification of the future stockpile. SNL also contributed to successful exhibits at SC10 and the second

Congressional Exhibition on Modeling and Simulation, as well as continued communications and strategic planning support to HQ.

In FY10, SNL organized and hosted the second Predictive Engineering Science Panel meeting; supported QASPR External Review Panel; supported PSAAP collaborations; provided support to the ASC Federal program office; developed and supported the SNL and HQ ASC Web sites; continued production of high-quality communications materials for HQ and the broader HPC community; supported the annual Principle Investigator's meetings; supported the Supercomputing Conference and Predictive Science Panel meetings; provided support for the ASC executive committee, support for quarterly meetings of the ASC executive committee, and management of the SAIC contract to provide various administration support to HQ; and supported collaborations with DTRA, STRATCOM, and the Nuclear Weapons Effects User's Group.

Planned activities in FY11:

- Host the Predictive Engineering Science Panel
- Support PSAAP collaborations
- Manage the SAIC contract to provide various administration support to HQ

Expected deliverables in FY11:

- Predictive Engineering Science Panel meeting
- Support for the SC10 Supercomputing Conference

Preliminary planned activities in FY12:

- Host PESP and support PSAAP
- Support programmatic needs of Exascale program
- Manage the SAIC contract to provide various administration support for HQ

WBS 1.5.5.4: System and Environment Administration and Operations

This level 4 product provides necessary operational support for reliable production computing and storage environments. The following activities are included: system administration and operations, software and hardware maintenance, licenses and contracts, computing environment security and infrastructure, requirements planning, initial deployment, production computing services, and tri-lab system integration and support.

System and Environment Administration and Operations Deliverables for FY11:

- Integration of Unclassified IDA cluster
- Muir system to GA status
- Integration of TLCC2 systems
- Ongoing support of Purple, the Sequoia ID (Dawn), BlueGene/L, Peloton, and TLCC capacity systems and visualization clusters
- Contracts and licenses needed for system operations and vendor support
- Self-maintenance of the Muir visualization cluster

- Ongoing, well-maintained Sequoia ID (Dawn), BlueGene/L, visualization clusters, parallel global file systems, Peloton, and TLCC capacity systems
- Self-maintenance of the Sequoia ID (Dawn)
- Deployment of identity management solution to replace existing LC Account Management System functionality
- Completed acceptance testing of Cielo system including Panasas file system
- Completed transition of Cielo into classified LANL network
- Completed Cielo system stabilization period and security test plans
- Single integrated security plan for all HPC clusters and storage servers
- Deployment of HPC monitoring infrastructure to all existing ASC platforms
- Self maintenance for Cielo, Roadrunner, and TLCC systems
- Integration of SCC environmental sensor software plus reports
- Ongoing support of ASC computing platforms and infrastructure
- 24x7 monitoring of ACS platforms and infrastructure
- Integration of TLCC2 systems and TOSS software stack
- Production ready TLCC platforms (in New Mexico and California)
- Production ready HPSS platforms
- Support of an encrypted file system for NSCC system
- Production ready mini-Cielo platform for National Security customers patterned after Cielo

WBS 1.5.5.4 System and Environment Administration and Operations (LLNL)

This product provides necessary operational support for reliable production computing environments. The following activities are included: system administration and operations, software and hardware maintenance, licenses and contracts, computing environment security and infrastructure, requirements planning, initial deployment, production computing services, and tri-lab system integration and support.

In FY10, LLNL integrated the Sequoia ID (Dawn) into GA status in the classified environment. The IDA cluster was brought into classified production status. Coastal (6 SUs) was moved from OCF to SCF in under six weeks. In Q4FY10, LLNL expects to take delivery of and integrate the Muir visualization system in anticipation of Sequoia. LLNL performed ongoing hardware self-maintenance of BlueGene/L, parallel global file systems, Peloton, and TLCC capacity systems. LLNL expanded its deployment of an identity management solution to support electronic workflows for approvals, management, and provisioning of groups. LLNL upgraded the security infrastructure to operate under the NAPS-compliant *LC Information System Security Plan*.

Planned activities in FY11:

- Decommission the Purple system
- Design, procure, and integrate the unclassified IDA cluster
- Install and integrate the Muir visualization cluster in the SCF

- Support the Sequoia ID (Dawn), Purple, BlueGene/L, Peloton, and TLCC capacity systems
- Re-engineer the inventory and RMA tracking system used for self-maintenance support
- Track and place contracts and licenses needed for system operations and vendor support
- Provide ongoing hardware self-maintenance of the Sequoia ID (Dawn), BlueGene/L, visualization clusters, parallel global file systems, Peloton, and TLCC capacity systems
- Continue upgrades to core security infrastructure components, thus enabling enhanced security capabilities and mechanisms to be offered and utilized
- Integrate security middleware software stack for latest TOSS release and planned TLCC systems
- Develop and integrate electronic workflow and provisioning functionality to replace the LC Account Management System, including annual revalidation and enterprise UID registry
- Participate in the effort to build a site security component library for applications and services, including preparing security test and evaluation plans
- Perform ongoing security-related activities in support of the secure use and secure management of ASC platforms and associated infrastructure

Expected deliverables in FY11:

- Integration of unclassified IDA cluster
- Muir system GA status
- Integration of TLCC2 systems
- Ongoing support of Purple, the Sequoia ID (Dawn), BlueGene/L, Peloton, and TLCC capacity systems and visualization clusters
- Delivery of first Sequoia racks
- Contracts and licenses needed for system operations and vendor support
- Self-maintenance of the Muir visualization cluster
- Ongoing, well-maintained Sequoia ID (Dawn), BlueGene/L, visualization clusters, parallel global file systems, Peloton, and TLCC capacity systems
- Self-maintenance of the Sequoia ID (Dawn)
- Deployment of identity management solution to replace existing LC Account Management System functionality
- Integration of the data management for RSA authentication manager with security registry and identity management services
- Deployment of data encryption capability for off-site disaster recovery backups
- Disaster recovery plan document and end-to-end recovery testing using off-site backups
- Deployment of a one-way link file interchange system to enable file transfers from unclassified to classified network

- Deployment of a standards-based one-time password authentication mechanism for the classified computing environment

Preliminary planned activities in FY12:

- Support the Sequoia ID (Dawn), TLCC1/TLCC2 capacity systems
- Integrate Sequoia in the classified environment
- Decommission of BlueGene/L and Peloton systems
- Bring TLCC2 systems into GA status
- Put in place contracts and licenses needed for system operations and vendor support
- Ongoing maintenance of BlueGene/L, the Sequoia ID (Dawn), parallel global file systems, Peloton, and TLCC capacity systems
- Begin self-maintenance of the Sequoia system
- Investigate issues and models for securely supporting and using a Cloud Computing paradigm
- Integrate the LC identity management solution with the LLNL Enterprise Identity Management deployment and identity aggregation efforts.
- Integrate security middleware software stack for next generation platforms

WBS 1.5.5.4 System Administration and Storage (LANL)

The System Administration and Storage project covers all services for computational systems operated by LANL for the purpose of providing an HPC production computing environment for weapons designers, developers, and engineers. The project works with users to troubleshoot problems experienced while running their applications, and helps users transition from old to new computing platforms. The capabilities include system configuration, system and user security, resource management, system administration and monitoring, archival storage, Panasas, and NFS file systems.

In FY10, LANL transitioned to limited availability of Roadrunner in support of ASC milestone work and deployed the initial system and application testbed systems for Cielo.

Planned activities in FY11:

- Support system by conducting ongoing and daily system and storage administration with continuous monitoring of production systems and infrastructure servers
- Ensure workload is carried out by proper configuration of queues and scheduling policies; daily monitoring and problem resolution of use problems
- Improve continuously the end-to-end level of service as seen by the users
- Conduct ongoing studies and improvement projects in the stability of large, integrated systems, including the development of improved diagnostic and monitoring capabilities
- Ensure data storage operations for GPFS, NFS, and archival storage (HPSS)
- Extend LANL and SNL best practices in system integration and support via the ACES partnership

Expected deliverables in FY11:

- Support for Cielo weapons science and select integrated weapons applications
- Contribution to the new releases of the CCE TOSS software stack
- Improved, HPC community leveraged monitoring infrastructure deployment to all existing ASC platforms
- Single, integrated security plan for all HPC clusters and storage servers
- Monitoring support for all ASC systems, including data storage

Preliminary planned activities in FY12:

- Transition Cielo to full tri-lab capability usage and support
- Support an expanded set of weapons applications on Cielo

WBS 1.5.5.4 Operations and Procurement Support (LANL)

The Operations and Procurement Support project provides around-the-clock operations and monitoring of the scientific computing resources, including performance computers such as the original Roadrunner Base System (Redtail and Yellowrail), Roadrunner Phase 3, Hurricane, Lobo, and data storage and retrieval systems such as the HPSS. In addition to monitoring all components 24x7, the computer operators provide systems hardware maintenance for all ASC platforms, including the new one petaFLOPS high-performance computer, Roadrunner Phase 3.

The Procurement Support aspect of this project assists customers with the technical and administrative aspects of planning, procurement, and contractual agreements for computer hardware and software products and services. This includes providing ongoing maintenance of computing hardware, software, and related support equipment, and analyst services through maintenance contracts.

This project utilizes the skill, experience, and talents of professional teams of operators and other skilled technical resources. These resources utilize automated tools and documented procedures to operate and monitor all HPC systems and components. This includes all components of the production computing environment, including compute engines, hardware, file servers, archival storage systems, the facilities they reside in and utilities they are dependent upon and all required software on these systems.

Planned activities in FY11:

- Provide around-the-clock operations and monitoring of the scientific computing resources, including an increased level of system hardware self-maintenance for various computing and data storage systems
- Ensure data storage operations for GPFS, NFS, and archival storage (HPSS)
- Continue to train operations staff in the hardware maintenance of Cielo and TLCC2
- Continue to work with system administrators to enhance the monitoring tool infrastructure to include additional features, such as asset tracking and ability to cascade to hierarchical windows for detecting node problems
- Continue to train operations staff in the use of the Zenoss monitoring tool for identifying hardware failures on all systems; assist in identifying and developing additional filters and functionality to enhance the effectiveness of Zenoss in

operations monitoring; expand the use of Zenoss to monitor the internal management network of all installed HPC systems

- Provide hardware support for installation and integration of additional scientific computing platforms, including Cielo and TLCC2
- Provide front-end system support by providing Dell hardware maintenance; continue to identify the necessary spare parts to maintain these servers
- Extend LANL and SNL best practices in systems operations, monitoring, and support via the ACES partnership
- Provide technical and administrative support in the procurement of computer hardware and software products and services for all ASC platforms

Expected deliverables in FY11:

- Operations staff trained in the hardware repair of Cielo, Dell, and TLCC2 systems
- Monitoring support for all ASC systems, including data storage
- Maintenance coverage for all ASC systems

Preliminary planned activities in FY12:

- Provide around-the-clock operations and monitoring of the scientific computing resources, including an increased level of system hardware self-maintenance for all current and future generations of ASC platforms and data storage systems
- Provide technical and administrative support in the procurement of ASC-related systems

WBS 1.5.5.4 Requirements Planning (LANL)

The Requirements Planning project covers the planning activities for the collection and statistical evaluation of user requirements for computing resources, development of new metrics, and data collection.

The primary capability of this project is to collect and understand user requirements for production computing resources and quality of service, and to develop new metrics, data collection, and analysis techniques to assist these purposes. The project will build on last year's work by completing the deployment of HPC monitoring software on remaining LANL clusters.

Planned activities in FY11:

- Design and develop software to process SCC facility sensor environmental data
- Design and develop asset tracking software
- Enhance current reporting capabilities

Expected deliverables in FY11:

- SCC environmental sensor software plus reports
- Deployment of monitoring software on remaining LANL clusters
- Asset tracking monitoring system software plus reports
- Enhanced reporting for system utilization and availability

Expected deliverables in FY12:

- Maintain and enhance monitoring infrastructure
- Integrate job and hardware performance data with system health information
- Deploy monitoring on new capacity clusters

WBS 1.5.5.4 Cielo Capability Computing Platform Initial Deployment (LANL)

The scope of the Cielo Capability Computing Platform Initial Deployment project is to take delivery and start the deployment of Cielo. This includes completing the acceptance tests, system integration into the LANL network, system stabilization, and transition into the classified network. The primary capabilities are acceptance and diagnostic testing, system stabilization, system integration into the yellow network, and transition to secure network.

Planned activities in FY11:

- Coordinate with Cray on the delivery of Cielo at LANL
- Develop acceptance test plan for Cielo
- Deploy initial Cielo and provide system integration
- Create plan for integrating Cielo into secure computing environment
- Develop integrated support structure with vendor for operational issues

Expected deliverables in FY11:

- Delivery of Cielo system
- Complete acceptance testing for Cielo
- Delivery of file system for Cielo
- Usage model for Cielo
- Evaluation of delivered system with diagnostic testing
- Cielo system integration into LANL network environment
- Transition of Cielo into classified LANL network
- Completion of system stabilization period
- Completion of security and security test plans

Preliminary planned activities in FY12:

- Provide production support

WBS 1.5.5.4 Production Computing Services (SNL)

The Production Computing Services project's goals are to operate and maintain all production platforms and associated support systems, and operate ASC capability and capacity platforms, data services and visualization systems, long-term hierarchical storage services, high-performance network systems, tri-lab compatible cyber authentication and authorization systems, and monitoring and reporting services. This project supports tri-lab capability platform resource allocations and coordinates with tri-

lab peers in establishing priority scheduling, if required. This project coordinates the integration and deployment of TLCC capacity systems into SNL's production computing environment, in collaboration with WBS 1.5.4.7, Common Computing Environment. Support of CCE common service and environment decisions and configuration management activities will also be provided.

This project has expertise in operating capacity computing clusters, integrating file servers at the system or facility-wide level, deploying new computing, storage and data management platforms, and in retiring end of life platforms. System administration for complex HPC environments is provided, as are design and development activities for new innovative computing platforms.

New capabilities added in FY10:

- Support for National Security computing needs
- Designate the existing facility as a National Security Computing Center supporting Focus Area 4 goals
- Deliver the initial National security mini-Cielo platform (designated Cielo del Sur)

Planned activities in FY11:

- Validation testing and site planning for TLCC2 systems due in FY11 / FY12
- Debugging connectivity or user environment issues associated with access by LLNL and SNL personnel to Cielo platform at LANL
- Begin decommissioning activities for end-of-life visualization platforms (Red Rose and Black Rose)
- Maintain and operate centralized Lustre file servers for unclassified and classified networks supporting TLCC1 systems
- Maintain and operate centralized long-term hierarchical file storage systems running HPSS (classified and unclassified)
- Demonstrate additional security capabilities within the National Security Computing Center to serve more than one customer group

Expected deliverables in FY11:

- Production ready TLCC platforms (in New Mexico and California)
- Production ready HPSS platforms
- Support of an encrypted file system for NSCC system
- Production ready mini-Cielo platform patterned after Cielo for National Security customers

Preliminary planned activities in FY12:

- Continue operations as above for remaining capacity platforms, storage services, and National Security systems

WBS 1.5.5.5: Facilities, Network, and Power

This level 4 product provides necessary physical facility and other utility infrastructure. The following activities are included: facilities infrastructure, classified and unclassified

facility networks, wide-area classified networks, ongoing network operations, infrastructure integration, and power.

Facilities, Network, and Power Deliverables for FY11:

- Design and installation of site preparation elements for Sequoia infrastructure installation, including the liquid cooling infrastructure, electrical distribution from bottom floor to the computer room, the electrical wiring and installation in the room itself, and mechanical and electrical infrastructure required to site Sequoia
- Final cooling tower cell installation
- Development of tools to perform power management of systems to prepare for exascale computing
- Upgraded DISCOM WAN (double the bandwidth)
- Redundant 10GE encryptor on DISCOM WAN
- OpenLabNet connection upgraded to 10GE
- Enhancements to electricity and cooling to support expected increase in computing capacity in the LDCC along with new system cooling requirements
- Completion of the electrical metering and wireless sensor projects
- A unified user support interface to all tri-lab users of Cielo and its associated development, test, and quality machines, appropriate to the level of availability of these machines
- Integration with other LANL and SNL support mechanisms to maintain consistent internal support structures

WBS 1.5.5.5 Facilities, Network and Power (LLNL)

The Facilities Network and Power project provides for the necessary physical facilities, utilities, and power capabilities to support staff and the ASC computing environment. Capabilities include adequate raised floor space, cooling facilities, and power to site large-scale ASC platforms. In addition, this project funds needed office, meeting room, and auxiliary space to enable a highly motivated and effective staff. Also included are classified and unclassified facility networks, wide-area classified networks, and ongoing network operations.

In FY10, LLNL completed the east room 7.5-MW electrical distribution expansion from the new electrical yard into the first level of the machine room and fully commissioned and completed activation testing of the electrical components from the first phase of construction of the B-453 15-MW power expansion (7.5MW west room electrical equipment). LLNL also completed and implemented energy savings initiatives resulting from self-benchmarking tools and computational fluid dynamic analysis for B-453 (the TSF). To date, the temperatures in the computer rooms and on the chilled water supply have been raised for an annual energy savings of over 47,500,000 kwh/yr. These energy saving initiatives assisted in B-453 achieving U.S. Green Building Council LEED Gold certification on December 24, 2009. LLNL deployed 10Gbit encryption on the DISCOM WAN.

Planned activities in FY11:

- Continue to track the progress of the institution project elements for increased redundancy and reliability of the lab electrical distribution system that will support the 15-MW electrical power expansion for B-453 for 2011 completion
- Commission and complete activation testing of the electrical components from the second phase of construction of the B-453 15-MW power expansion (7.5MW east room electrical equipment)
- Maintain and support equipment in existing computational and staff facilities
- Continue analysis of future modifications and /or expansion of facilities that will be needed by future ASC systems
- Continue to implement the results from self-benchmarking tool for all of the computer rooms in B-453 (the TSF), B-451, B-439, B-115, and B-117 created by DOE Office of Science and Lawrence Berkeley National Laboratory (<http://hightech.lbl.gov/datacenters.html>) to continue to routinely identify prospective energy savings initiatives as computer rooms change
- Develop design concepts to perform the implementation of free cooling to improve PUE from 1.34 to 1.2
- Continue to update computational fluid dynamics model for all ASC systems to profile the airflow required to cool the machines adequately as platforms are added and retired
- Continue ongoing support of facility networks
- Continue evaluation of 40GE and 100GE technologies
- Continue evaluation of QDR and EDR Infiniband
- Re-bid DISCOM WAN contract
- Implement network for TLCC2 systems
- Deploy 10GE OpenLabNet connection

Expected deliverables in FY11:

- Design and installation of site preparation elements for Sequoia infrastructure installation. This includes the liquid cooling infrastructure, electrical distribution from bottom floor to the computer room, the electrical wiring and installation in the room itself and mechanical and electrical infrastructure required to site Sequoia. The end product of this milestone is electrical and mechanical capabilities available for successful siting of the initial delivery of Sequoia.
- Final cooling tower cell installation.
- Tools to perform power management of systems to prepare for exascale computing.
- Upgraded DISCOM WAN (double the bandwidth).
- Implementation of redundant 10GE encryptor on DISCOM WAN.
- OpenLabNet connection upgraded to 10GE.

Preliminary planned activities in FY12:

- Install final four air handlers in the mechanical rooms
- Continue to maintain and support the equipment in existing computational and staff facilities

- Implement tools to perform power management of systems to prepare for exascale computing
- Provide ongoing support of facility networks and DISCOM WAN

WBS 1.5.5.5 Facilities and Power (LANL)

The Facilities and Power project is responsible for the engineering, design, operation, and maintenance of the mission-important electrical, mechanical, cooling, and other computing infrastructure in support of the ASC program. The project provides support for infrastructure design upgrades, project management, user interface and oversight, demolition and decommissioning of older systems, and computer site preparation for new platforms. Because the tri-lab community requires the systems to be operational at all times, the project provides on-call support after hours and on weekends for facility related issues. This project covers the cost of electrical power in the SCC required by the high performance computers and associated cooling requirements.

This is a new project for FY11.

Planned activities in FY11:

- Continue operations and maintenance of electrical and mechanical systems for ASC computing facilities
- Provide facility support for decommissioning of the View Master system
- Site prep new View Master at SCC and View Master light show at LDCC
- Increase PUE and data center efficiency by implementing engineering strategies in an effort to optimize energy efficiencies and cost savings in computing facilities
- Continue involvement in the power and cooling design requirements for Cielo and future platforms
- Continue analysis and solutions to reduce PUE in LDCC and SCC
- Work closely with lab subject matter experts in planning future electrical power and cooling needs for exascale computing

Expected deliverables in FY11:

- Enhancements to electricity and cooling to support expected increase in computing capacity in the LDCC, along with new system cooling requirements
- Completed electrical metering and wireless sensor projects, which will help to identify power usage as well as monitor supply and exhaust temperatures and under floor static pressure

Preliminary planned activities in FY12:

- Continue to work with lab subject matter experts to address power and cooling design requirements for “exascale” computing
- Continue to focus on the facility’s site preparation requirements (including power, cooling, and space) for current and future HPC installations
- Continue operations and maintenance of electrical and mechanical systems for ongoing programmatic computing, providing the facility support for decommissioning of retired systems

- Increase PUE and data center efficiency in the computing facilities by implementing engineering strategies in an effort to optimize energy efficiencies and cost savings in computing facilities.
- Develop a long-term project planning tool for future of ASC computing projections with input from power and cooling subject matter experts and other supercomputing expertise

WBS 1.5.5.5 Networking (LANL)

The Networking project includes all services for networks operated by LANL for the purpose of providing an HPC networking environment for weapons designers, developers, and engineers and provides ongoing network operations that support ASC computing in the classified and unclassified networks. This includes directly attached networks to HPC systems (machine area network), network backbones, user LANs, and the high-end DisCom WAN connecting the tri-labs. Core capabilities include designing, developing, deploying, and supporting classified and unclassified network hardware and services to support ASC computational systems and infrastructure.

In FY10, LANL implemented a full 12-lane infrastructure between all classified ASC supercomputing systems and associated file systems. Numerous network infrastructure upgrades were completed to provide better capability via newer hardware, in addition to the addition of increased fiber trunk count between buildings at the LANL site. Support for HPC networking was more closely integrated into the mainstream HPC support capability.

Planned activities in FY11:

- Operate and maintain network services
- Support InfiniBand and Cielo proprietary torus interconnect fabrics
- Manage the high-performance network backbone
- Refine network to improve performance, reliability, and availability of backbone and services
- Increase reliability for HPC cluster management switches and core switches
- Upgrade network infrastructure, where necessary, to increase bandwidth, reliability, and availability to HPSS, visualization platforms, and customer workstations
- Provide support for system interconnects monitoring, reliability, and performance management

Expected deliverables in FY11:

- Full integration of high-level technical networking support into HPC division
- New technology deployment to increase network utilization, bandwidth, and reliability
- Ongoing network support for all ASC systems
- Integrate IHPC hardware, if required by IHPC project
- Integrate Cielo into existing network infrastructure
- Extend Zenoss HPC monitoring infrastructure to HPC cluster management network switches and core switches for increased reliability

Preliminary planned activities in FY12:

- Continue to operate and maintain LAN/Metropolitan Area Network/WAN infrastructure
- Plan I/O infrastructure to support next cluster platform

WBS 1.5.5.5 Facilities, Networking, and Power (SNL)

This SNL Project funds the power and space charges assigned to HPC systems (capacity and file system servers) and long-term hierarchical storage servers (running the HPSS software product). It provides for facilities and personnel to manage installation and removal of computing platforms, file systems, visualization systems, networking equipment, power distribution systems, and cooling systems in support of all computing resources. It also funds major operations contracts such as the ASC DISCOM WAN.

Facilities professionals have reduced the overall operating expenses by minimizing cooling and electrical distribution expenses over the last several years through a comprehensive program of introducing more efficient computer room air conditioning units, using higher voltage electrical source power distribution units, exploring alternative energy sources and conservation mechanisms, which include reducing the volume of chilled water required for cooling and improving air flow in the facility by minimizing obstructions underneath the computer floor.

In FY10, SNL realized significant cost savings in power and water by containment of air cooled systems. The tri-lab capacity computing platforms (Glory, Unity) permitted raising of cooling water by 5 degrees Fahrenheit, thus extending the period when chilled plate cooling can replace active chilling via the refrigerant cooler facilities plant system.

Planned activities in FY11:

- Prepare facility for TLCC2 system installs; plan for facility upgrades/modifications as necessary once design is known
- Install TLCC2 system as needed
- Complete installation of mini-Cielo platform in NSCC

Expected deliverables in FY11:

- Payment of power and space charges for all operational ASC systems
- Plans for installation of future TLCC2 platforms

Preliminary planned activities in FY12:

No facilities modifications or upgrades are planned at this time.

V. ASC Level 1 and 2 Milestones

Table V-1. ASC Level 1 *Proposed* Milestones and Interfaces with Defense Programs Components from FY12–FY16

Milestone Title	Level	FY	Completion Date	Site(s)	Participating Program Offices
Develop, implement, and apply a suite of physics-based models and high-fidelity databases necessary for predictive simulation of the initial conditions for primary boost (initial conditions 1)	1	FY12	Q4	LANL, LLNL	Science Campaigns ASC Campaign
Assessment of weapon surety status	1	FY13	TBD	SNL	ASC Campaign Engineering Campaigns
Demonstrate predictive capability for weapon system response to short-pulsed neutrons in hostile radiation environment	1	FY13	TBD	SNL	ASC Campaign
Baseline demonstration of UQ aggregation methodology for full-system weapon performance prediction	1	FY14	TBD	LANL, LLNL, SNL	Science Campaigns ASC Campaign DSW
Full-system safety assessment	1	FY14	TBD	SNL	ASC Campaign Engineering Campaigns
Advanced models to support initial conditions for boost (initial conditions 2)	1	FY14	TBD	LANL, LLNL	Science Campaigns ASC Campaign

Table V-2. Quick Look: Level 2 Milestone Dependencies for FY11⁵

Milestone ID	Milestone Title	Level	FY	Completion Date	DOE Program/Subprogram(s)	Site(s)
TBD	Improve physics and geometric fidelity simulation capability in nuclear performance code system	2	FY11	Mar-2011	IC	LLNL
TBD	Demonstrate progress on delivering aboveground experiment simulation capabilities	2	FY11	Jun-2011	IC	LLNL
TBD	Validate thermo-mechanical coupling	2	FY11	Sep-2011	IC	LLNL
TBD	Multiphase plutonium equation of state variations	2	FY11	Sep-2011	PEM	LLNL
TBD	First principles equation of state of high-explosives binders	2	FY11	Jun-2011	PEM	LLNL
TBD	Numerical convergence study of selected burn calculations	2	FY11	Mar-2011	V&V	LLNL
TBD	Primary Metrics Project event expanded to over 50 events	2	FY11	Mar-2011	V&V	LLNL
TBD	Uncertainty quantification study of Secondary Computational Assessment and Metrics Project (SCAMP) survey events	2	FY11	Sep-2011	V&V	LLNL
TBD	Scalable applications preparation and outreach for Sequoia	2	FY11	Jun-2011	CSSE	LLNL
TBD	Deploy high-performance storage system quota system	2	FY11	Dec-2010	CSSE	LLNL
TBD	Early users on LLNL TLCC2	2	FY11	Sep-2011	FOUS	LLNL
TBD	Sequoia facilities integration	2	FY11	Sep-2011	FOUS	LLNL
TBD	Develop a Lagrange application design capability to support Directed Stockpile Work and the National Code Strategy	2	FY11	Jun-2011	IC	LANL
TBD	Enhance the Eulerian applications algorithms and performance in support of the National Code Strategy	2	FY11	Jun-2011	IC	LANL
TBD	Next generation of charged-particle data capabilities	2	FY11	Sep-2011	PEM	LANL
TBD	Initial implementation of in-line non-local thermodynamic equilibrium capability	2	FY11	Jun-2011	PEM	LANL
TBD	Thermonuclear applications verification and validation assessment of physics modeling capabilities in an ASC code near thresholds	2	FY11	Sep-2011	V&V	LANL
TBD	Data restructuring in a production application	2	FY11	Sep-2011	CSSE	LANL

⁵ Factors such as FY11 Congressional Appropriations, NNSA/DP directives, and National Security considerations may necessitate a change in the current milestone set.

Milestone ID	Milestone Title	Level	FY	Completion Date	DOE Program/Subprogram(s)	Site(s)
TBD	Time-dependent NWM21 radiation environment	2	FY11	Mar-2011	IC	SNL
TBD	III-V heterojunction transistor models for hostile environments	2	FY11	Sep-2011	IC	SNL
TBD	Blast and fragmentation abnormal environment modeling for the B61	2	FY11	Sep-2011	IC	SNL
TBD	Models for high voltage breakdown in neutron tubes	2	FY11	Sep-2011	PEM	SNL
TBD	Ductile failure X-Prize	2	FY11	Sep-2011	PEM	SNL
TBD	Cavity system-generated electromagnetic pulse verification and validation	2	FY11	Dec-2010	V&V	SNL
TBD	Develop feedback system for intelligent dynamic resource allocation to improve application performance	2	FY11	Sep-2011	CSSE	SNL
TBD	Cielo capability computing platform integration readiness	2	FY11	Dec-2010	CSSE	LANL, SNL
TBD	TLCC2 contract awarded	2	FY11	Mar-2011	CSSE	LLNL, LANL, SNL
TBD	Deployment of a common capacity computing environment	2	FY11	Sep-2011	CSSE	LLNL, LANL, SNL

Table V-3. Quick Look: *Preliminary* Level 2 Milestone Dependencies for FY12

Milestone ID	Milestone Title	Level	FY	Completion Date	DOE Program/Subprogram(s)	Site(s)
TBD	Implement improved algorithms for loading on structures	2	FY12	Mar-2012	IC	LLNL
TBD	Extend ALE AMR capabilities to additional physics packages	2	FY12	Sep-2012	IC	LLNL
TBD	Low-pressure multiscale strength model	2	FY12	Sep-2012	PEM	LLNL
TBD	Continued uncertainty quantification study of Secondary Computational Assessment and Metrics Project (SCAMP) survey events	2	FY12	Sep-2012	V&V	LLNL
TBD	Scalable applications preparations and outreach for Sequoia	2	FY12	Jun-2012	CSSE	LLNL
TBD	Deploy SLURM on Sequoia	2	FY12	Sep-2012	CSSE	LLNL
TBD	Improve the Eulerian applications codes through additional physics and algorithms to allow validation via high energy density experiments	2	FY12	Mar-2012	IC	LANL
TBD	Expand Lagrange application 3D hydro capabilities in support of Directed Stockpile Work	2	FY12	Jun-2012	IC	LANL
TBD	Production release of the neutron transport Monte Carlo application toolkit	2	FY12	Jun-2012	IC	LANL
TBD	New constitutive thermo-mechanical model for high explosives	2	FY12	Dec-2011	PEM	LANL
TBD	Baseline an ASC code for an emergency response application and develop plan to meet future requirements	2	FY12	Mar-2012	PEM	LANL
TBD	SESAME uranium update	2	FY12	Mar-2012	PEM	LANL
TBD	Common mix modeling	2	FY12	Sep-2012	PEM	LANL
TBD	Engineering verification and validation assessment of a weapon subassembly in an ACRR experiment	2	FY12	Sep-2012	V&V	LANL
TBD	Assessment of predictive capability for primary performance using LANL primary validation suite	2	FY12	Mar-2012	V&V	LANL
TBD	Benchmark evaluation of predictive capability for boost using LANL boost validation suite	2	FY12	Jun-2012	V&V	LANL
TBD	Programming models and data analysis environments for extreme-scale computing	2	FY12	Sep-2012	CSSE	LANL

Milestone ID	Milestone Title	Level	FY	Completion Date	DOE Program/Subprogram(s)	Site(s)
TBD	Application deployment of a quick parallel log-structured file system capability	2	FY12	Sep-2012	CSSE	LANL
TBD	High-performance storage system 8.1 deployment	2	FY12	Sep-2012	CSSE	LANL
TBD	Visualization cluster upgrade project	2	FY12	Dec-2011	CSSE	LANL
TBD	Cavity SGEMP predictive capability for realistic RB geometry	2	FY12	Mar-2012	IC	SNL
TBD	One-way coupling of re-entry aerodynamics with ablation to structural analysis simulations	2	FY12	Sep-2012	IC	SNL
TBD	Sierra Toolkit conversion, Phase 2	2	FY12	Sep-2012	IC	SNL
TBD	Gas transfer system performance models for hydrogen embrittlement	2	FY12	Sep-2012	PEM	SNL
TBD	B61 normal environment tape joint performance	2	FY12	Sep-2012	PEM	SNL
TBD	Computational uncertainty quantification for the QASPR Project	2	FY12	Sep-2012	V&V	SNL
TBD	Enhanced solution verification capability in Encore, applied to abnormal thermal quantification of margins and uncertainties	2	FY12	Sep-2012	V&V	SNL
TBD	Verification and validation of normal mechanical environments for B61 System	2	FY12	Sep-2012	V&V	SNL
TBD	Demonstration of a legacy application's path to exascale	2	FY12	Mar-2012	CSSE	SNL
TBD	Cielo capability computing platform production readiness	2	FY12	Dec-2011	CSSE	LLNL, LANL, SNL
TBD	Deploy TLCC2 and a common capacity computing environment	2	FY12	Sep-2012	CSSE	LLNL, LANL, SNL

Detailed Milestone Descriptions for FY11

Milestone (ID#): Scalable applications preparation and outreach for Sequoia		
Level: 2	Fiscal Year: FY11	DOE Area/Campaign: ASC
Completion Date: June 30, 2011		
ASC nWBS Subprogram: CSSE		
Participating Sites: LLNL		
Participating Programs/Campaigns: ASC		
Description: Building on the FY10 SAP milestone, the SAP effort will extend the knowledge base, documentation, and training to enable ASC code teams to utilize Sequoia. SAP will actively engage tri-lab code teams to address their needs in porting codes to the Sequoia ID (Dawn) and in preparing for the arrival of Sequoia in FY12. For FY11, several multi-physics codes will be engaged to characterize the Sequoia ID (Dawn) performance, analyze bottlenecks and load balance issues, and to develop strategies for improving performance targeting the Sequoia system. Additional refinements to understanding and performance enhancements will be achieved for codes targeted in the initial SAP effort. A deployment plan for development environment tools will be developed, including initial testing and validation of the tool suite, using Sequoia simulators.		
Completion Criteria: A report covering the performance findings and recommended techniques and strategies for the codes studied.		
Customer: ASC IC		
Milestone Certification Method: Professional documentation, such as a report or a set of viewgraphs with a written summary, is prepared as a record of milestone completion. The "handoff" of the developed capability (product) to a nuclear weapons stockpile customer is documented.		
Supporting Resources: ASC the Sequoia ID (Dawn) system		
Supporting Milestones:		
Program	Title	Due Date
N/A	N/A	N/A
Codes/Simulation Tools Employed: The Sequoia ID (Dawn) software environment and IBM Sequoia simulators		
Contribution to the ASC Program: More effective use of Sequoia computer system		
Contribution to Stockpile Stewardship: More effective use of Sequoia computer system		

No.	Risk Description	Risk Assessment (low, medium, high)		
		Consequence	Likelihood	Exposure
	No access to prototype Sequoia hardware	Low	Moderate	Low

Milestone (ID#): Deploy high-performance storage system quota system				
Level: 2	Fiscal Year: FY11	DOE Area/Campaign: ASC		
Completion Date: December 30, 2010				
ASC nWBS Subprogram: CSSE				
Participating Sites: LLNL				
Participating Programs/Campaigns: ASC				
Description: The LC's HPSS archive annual growth will soon exceed projected supporting budgets. The Annual Archive Allowance System (AAAS) will allow LC management and administrators to establish annual storage growth allowances (quotas) in partnership with users, programs, and projects. Users, managers, and administrators will have role-based means by which to monitor and manage archival storage growth both on a per-user and project basis. When allowances are exceeded, notifications will be made to users and warnings will be posted in management-viewable reports.				
Completion Criteria: This project is complete when the AAAS monitoring and administration tool is deployed to LC users in OCF and SCF environments.				
Customer: LC				
Milestone Certification Method: A report will be written that documents deployment of the AAAS system and infrastructure in the LC A memo will be written that documents the handoff of the developed and deployed capability to Michel McCoy, ASC Program Leader at LLNL.				
Supporting Resources: LC High Performance Systems Division, CSSE personnel				
Supporting Milestones:				
	Program	Title	Due Date	
	N/A	N/A	N/A	
Codes/Simulation Tools Employed: None				
Contribution to the ASC Program: Required for successful management of ASC archival storage system resources in order to allow simulation results to continue being cost-effectively retained within LC archives.				
Contribution to Stockpile Stewardship: This is a tool that will help reduce costs associated with archival storage of stockpile stewardship data by highlighting and managing storage use.				
No.	Risk Description	Risk Assessment (low, medium, high)		
		Consequence	Likelihood	Exposure
1	AAAS is not integrated in a timely fashion, users exceed budgeted archival storage space resulting in the inability to store subsequent simulation results in the archive	High	Low	Low

Milestone (ID#): Early users on LLNL TLCC2		
Level: 2	Fiscal Year: FY11	DOE Area/Campaign: ASC
Completion Date: September 30, 2011		
ASC nWBS Subprogram: FOUS		
Participating Sites: LLNL		
Participating Programs/Campaigns: ASC		
Description: This milestone encompasses all the work past the TLCC2 contract award associated with the TLCC2 system integration, including hardware and software troubleshooting, required electrical and mechanical facilities integration necessary for system deployment, acceptance testing of hardware, and deployment and integration of at least one SU on the unclassified network. The end product of this milestone will be one or more Limited Availability TLCC2 SUs functioning properly on the unclassified network with at least one early user running science codes.		
Completion Criteria: This project is complete when a user is running science codes on one or more TLCC2 SUs in the LC's OCF environment.		
Customer: ASC		
Milestone Certification Method: A report will be written that documents completion of the synthetic workload testing, as well as acceptance and deployment of a TLCC2 SU to the unclassified network in the LC. A memo will be written that documents handoff of the developed and deployed system to Michel McCoy, ASC Program Leader at LLNL.		
Supporting Resources: LC High Performance Systems Division, CSSE personnel		
Supporting Milestones:		
Program	Title	Due Date
N/A	N/A	N/A
Codes/Simulation Tools Employed: N/A		
Contribution to the ASC Program: Required to provide capacity computing cycles necessary for simulation in support of stockpile stewardship.		
Contribution to Stockpile Stewardship: Provides capacity computing cycles necessary for stockpile stewardship.		

No.	Risk Description	Risk Assessment (low, medium, high)		
		Consequence	Likelihood	Exposure
1	Hardware issues in TLCC2 delay ability to deploy	High	Moderate	Medium
2	Technology production issues from vendors delay ability to deploy TLCC2	High	Moderate	Medium
3	Delays in the TLCC2 contract award process impacts ability to deploy	High	Low	Medium

Milestone (ID#): Sequoia facilities integration				
Level: 2	Fiscal Year: FY11	DOE Area/Campaign: ASC		
Completion Date: September 30, 2011				
ASC nWBS Subprogram: FOUS				
Participating Sites: LLNL				
Participating Programs/Campaigns: ASC				
Description: This milestone encompasses the work required to site the Sequoia system. This includes the liquid cooling infrastructure, electrical distribution from bottom floor to the computer room, the electrical wiring and installation in the room itself, and mechanical and electrical infrastructure required to site Sequoia. The end products of this milestone are electrical and mechanical capabilities available for successful siting of the initial delivery of Sequoia.				
Completion Criteria: This project is complete when a memo certifying Sequoia facility readiness is delivered to Michel McCoy, ASC Program Leader.				
Customer: ASC				
Milestone Certification Method: A report will be written that documents completion of the mechanical and electrical facilities necessary to site Sequoia. A memo addressed to Michel McCoy, ASC Program Leader at LLNL, will be written to document the readiness of the facility for Sequoia siting.				
Supporting Resources: LC High Performance Systems Division, CSSE personnel				
Supporting Milestones:				
Program	Title		Due Date	
N/A	N/A		N/A	
Codes/Simulation Tools Employed: N/A				
Contribution to the ASC Program: Required to enable deployment of Sequoia computing cycles necessary for simulation (UQ and Boost) in support of stockpile stewardship.				
Contribution to Stockpile Stewardship: Provides computing cycles necessary for stockpile stewardship.				
No.	Risk Description	Risk Assessment (low, medium, high)		
		Consequence	Likelihood	Exposure
1	Modifications in vendor scope for the solution.	High	Moderate	High
2	Delay in design/build construction completion of highly complex systems.	High	Moderate	High

3	Commissioning first machine level liquid cooling system from vendor.	Moderate	Moderate	High
4	Uncertainty in projected electrical consumption for the full computer.	High	High	Medium

Milestone (ID#): Data Restructuring in a Production Application				
Level: 2	Fiscal Year: FY11	DOE Area/Campaign: ASC		
Completion Date: September 30, 2011				
ASC nWBS Subprogram: CSSE				
Participating Sites: LANL				
Participating Programs/Campaigns: ASC				
Description: Implement advanced data structures, communications, and parallel I/O for visualization in ASC codes through ASC Eulerian refactoring project, and demonstrate the efficiency and improvement of the new implementations through practical problems. This milestone will lead to the subsequent use of the advanced data structures in applications.				
Completion Criteria: Results documented as a report. Source codes the advanced data structures, communications, and parallel I/O committed to the refactored ASC codes.				
Customer: ASC IC				
Milestone Certification Method: A program review is conducted and its results are documented. Professional documentation, such as a report or a set of viewgraphs with a written summary, is prepared as a record of milestone completion.				
Supporting Resources: Access to supercomputer resources				
Supporting Milestones:				
	Program	Title	Due Date	
	ASC/CSSE	Cielo capability computing platform production capability readiness	Jun-11	
	ASC/IC	ASC Eulerian codes	Sep-11	
Codes/Simulation Tools Employed: ASC code project				
Contribution to the ASC Program: The milestone will provide the ASC community with the new parallel capability of the refactored Eulerian codes.				
Contribution to Stockpile Stewardship: The refactored Eulerian code will provide new physics capability, which Eulerian codes previously did not provide.				
No.	Risk Description	Risk Assessment (low, medium, high)		
		Consequence	Likelihood	Exposure
1	Loss of key personnel	High	Moderate	Medium

Milestone (ID#): Develop feedback system for intelligent dynamic resource allocation to improve application performance								
Level: 2	Fiscal Year: FY11	DOE Area/Campaign: ASC						
Completion Date: September 30, 2011								
ASC nWBS Subprogram: CSSE								
Participating Sites: SNL								
Participating Programs/Campaigns: ASC								
Description: Demonstrate the ability to affect application performance on a Cielo-like architecture or similar-scale platform through appropriate resource allocation (both static and dynamic) based on historic, run-time, or user-furnished process resource requirements in conjunction with current platform resource usage and availability. Historic and run-time application resource requirement characteristics as well as platform resource utilization will be derived using a scalable information gathering and run-time analysis system. Numeric hardware related metrics and information acquired from the platform scheduler/resource manager will be used.								
Completion Criteria: Completion of program review and final document published as a SAND report.								
Customer: Platform design team members and algorithm/application developers.								
Milestone Certification Method: A program review is conducted and its results are documented. Professional documentation, such as a report or a set of viewgraphs with a written summary, is prepared as a record of milestone completion.								
Supporting Resources: CSSE Simulation & Tools program, CSSE Advanced Systems program								
Supporting Milestones: <table> <tr> <th>Program</th><th>Title</th><th>Due Date</th></tr> <tr> <td>N/A</td><td>N/A</td><td>N/A</td></tr> </table>			Program	Title	Due Date	N/A	N/A	N/A
Program	Title	Due Date						
N/A	N/A	N/A						
Codes/Simulation Tools Employed: OVIS, HERMES								
Contribution to the ASC Program: Enable capabilities for improved ASC system utilization and ASC application performance.								
Contribution to Stockpile Stewardship: Enable capabilities for improved application performance, particularly in production system scenarios.								

No.	Risk Description	Risk Assessment (low, medium, high)		
		Consequence	Likelihood	Exposure
1	Lack of access to large-scale platform	Low	Moderate	Low
2	Lack of access to resource state information via kernel-based information gathering system	Moderate	High	Medium
3	Lack of a system administrator approval to incorporate data collection	Moderate	Low	Medium

Milestone (ID#): Cielo capability computing platform integration readiness				
Level: 2		Fiscal Year: FY11		DOE Area/Campaign: ASC
Completion Date: December 31, 2010				
ASC nWBS Subprogram: CSSE				
Participating Sites: LANL, SNL				
Participating Programs/Campaigns: ASC				
Description: Ready Cielo for integration into the LANL computing center. Deliver and install system hardware. Deliver, test, and demonstrate system software. Complete on-site capability scaling testing. Ready Cielo for on-site integration into the local and remote computing infrastructure, including the user software environment.				
Completion Criteria: Follows the ASC Level 2 Milestone criteria for capability platforms: system hardware deliveries from vendor to site are complete, including the basic hardware to integrate "the system" as contractually defined; installation of the system by the contractor on-site to the extent that is contractually required is substantially complete; in general, contractual requirements for formal hardware acceptance have been substantially completed; system software needed for basic operation of the system is delivered, tested, and demonstrated to be operational; vendor has completed on-site capability scaling testing and demonstration; and system is ready to begin on-site integration into local computing environment.				
Customer: NNSA / ASC HQ, tri-lab ASC program managers responsible for CCCs, SSP, tri-lab weapons applications community.				
Milestone Certification Method: A program review is conducted and its results are documented. Professional documentation, such as a report or a set of viewgraphs with a written summary, is prepared as a record of milestone completion.				
Supporting Resources: CSSE, FOUS, platform funding, ACES program managers, LANL facilities.				
Supporting Milestones:				
Program		Title		Due Date
N/A		N/A		N/A
Codes/Simulation Tools Employed: N/A				
Contribution to the ASC Program: Provides production capability compute cycles to ASC Program including scalable performance.				
Contribution to Stockpile Stewardship: Primary production capability platform for the ASC Program.				
No.	Risk Description	Risk Assessment (low, medium, high)		
		Consequence	Likelihood	Exposure
1.	Contract not awarded on schedule	High	Low	Medium

Milestone (ID#): TLCC2 contract awarded				
Level: 2	Fiscal Year: FY11	DOE Area/Campaign: ASC		
Completion Date: March 31, 2011				
ASC nWBS Subprogram: CSSE				
Participating Sites: LLNL, LANL, SNL				
Participating Programs/Campaigns: ASC				
Description: Develop, issue, and evaluate the RFP for the TLCC2 platform and award a contract.				
Completion Criteria: Contract awarded to winning vendor for TLCC procurement				
Customer: ASC HQ				
Milestone Certification Method: A program review is conducted and its results are documented. Professional documentation, such as a report or a set of viewgraphs with a written summary, is prepared as a record of milestone completion.				
Supporting Resources: CSSE, FOUS, procurement staff				
Supporting Milestones:				
Program		Title		Due Date
N/A		N/A		N/A
Codes/Simulation Tools Employed: N/A				
Contribution to the ASC Program: Common platforms leverage resources and enhance user efficiency.				
Contribution to Stockpile Stewardship: Provides capacity computing cycles necessary for stockpile stewardship.				
No.	Risk Description	Risk Assessment (low, medium, high)		
		Consequence	Likelihood	Exposure
1	COTS vendor silicon non-functional	High	Low	Low
2	Integrator unable to perform	High	Low	Low

Milestone (ID#): Deployment of a common capacity computing environment				
Level: 2	Fiscal Year: FY11	DOE Area/Campaign: ASC		
Completion Date: September 30, 2011				
ASC nWBS Subprogram: CSSE				
Participating Sites: LLNL, LANL, SNL				
Participating Programs/Campaigns: ASC				
Description: Deploy additional CCE capabilities for capacity computing environment, working towards a responsive and more efficient infrastructure to support computing for QMU and predictivity.				
Completion Criteria: Deploy CCE capabilities developed during FY10 including: major upgrades to the common operating system (TOSS 2.0); OI SS; Workload Characterization; and Gazebo Test and Analysis Suite; and new capabilities in the Debugger, Security Integration and OpenMPI projects. Demonstrate the tri-lab CCE software stack on the production ASC TLCC systems. Develop new capabilities in the continuing FY11 CCE projects. Prepare for deployment of the next generation of the ASC TLCC systems, which may include hardware and software integration and testing for the tri-lab environment. The tri-labs will continue to do gap and risk analysis of the CCE software stack and add new projects, as needed, to address high-priority gaps.				
Customer: Users of the ASC capacity systems				
Milestone Certification Method: A program review is conducted and its results are documented. Professional documentation, such as a report or a set of viewgraphs with a written summary, is prepared as a record of milestone completion.				
Supporting Resources: ASC CSSE sub-program				
Supporting Milestones:				
	Program	Title	Due Date	
	N/A	N/A	N/A	
Codes/Simulation Tools Employed: N/A				
Contribution to the ASC Program: CCE for capacity computing				
Contribution to Stockpile Stewardship: Easy-to-use capacity computing resource				
No.	Risk Description	Risk Assessment (low, medium, high)		
		Consequence	Likelihood	Exposure
	none			

Milestone Descriptions for Preliminary FY12

Milestone (ID#): Scalable applications preparations and outreach for Sequoia		
Level: 2	Fiscal Year: FY12	DOE Area/Campaign: ASC
Completion Date: June 30, 2012		
ASC nWBS Subprogram: CSSE		
Participating Sites: LLNL		
Participating Programs/Campaigns: ASC		
Description: With the delivery of the Sequoia system in 1QFY12, the SAP effort will fully deploy the applications development environment for the new system. SAP will evaluate the utility of the environment, address deficiencies that may be identified, and assist initial applications code teams with issues encountered. Preliminary performance of applications and benchmarks on the Sequoia hardware will be evaluated.		

Milestone (ID#): Deploy SLURM on Sequoia		
Level: 2	Fiscal Year: FY12	DOE Area/Campaign: ASC
Completion Date: September 30, 2012		
ASC nWBS Subprogram: CSSE		
Participating Sites: LLNL		
Participating Programs/Campaigns: ASC		
Description: The SLURM resource manager will be ported to Sequoia.		

Milestone (ID#): Programming models and data analysis environments for extreme-scale computing		
Level: 2	Fiscal Year: FY12	DOE Area/Campaign: ASC
Completion Date: September 30, 2012		
ASC nWBS Subprogram: CSSE		
Participating Sites: LANL		
Participating Programs/Campaigns: ASC/CSSE		
Description: Explore, quantitatively investigate, and report on the technological choices for programming models and data analysis environments on petascale, potentially hybrid, multi-core computing systems. Establish plans and goals for FY13 and FY14 for enabling application development for extreme-scale computer systems.		

Milestone (ID#): Application deployment of a quick parallel log-structured file system capability		
Level: 2	Fiscal Year: FY12	DOE Area/Campaign: ASC
Completion Date: September 30, 2012		
ASC nWBS Subprogram: CSSE		
Participating Sites: LANL		
Participating Programs/Campaigns: ASC/CSSE		
Description: Deployment and assessment of PLFS with HASH in multiple production ASC applications on a production machine.		

Milestone (ID#): High-performance storage system 8.1 deployment		
Level: 2	Fiscal Year: FY12	DOE Area/Campaign: ASC
Completion Date: September 30, 2012		
ASC nWBS Subprogram: CSSE		
Participating Sites: LANL		
Participating Programs/Campaigns: ASC/CSSE		
Description: Includes 8.1 placed in GA in 2nd quarter of FY12. Installed in LANL open environment in 3rd quarter FY12. Installed in LANL secure environment 4th quarter FY12.		

Milestone (ID#): Visualization cluster upgrade project		
Level: 2	Fiscal Year: FY12	DOE Area/Campaign: ASC
Completion Date: December 30, 2011		
ASC nWBS Subprogram: CSSE		
Participating Sites: LANL		
Participating Programs/Campaigns: ASC		
Description: The focus of this milestone will be to provide the necessary resources to visualize output generated on petascale clusters. Requirements will be developed, equipment will be purchased, and the cluster will be integrated into the LANL computer center. The milestone will be complete when the cluster is ready for production work.		
Completion Criteria: Results documented as a report. System made available to the ASC community.		
Customer: Weapons user community		

Milestone (ID #): Visualization cluster upgrade project				
Milestone Certification Method: A program review is conducted and its results are documented. Professional documentation, such as a report or a set of viewgraphs with a written summary, is prepared as a record of milestone completion.				
Supporting Resources: Access to Roadrunner and other supercomputer resources				
Supporting Milestones:				
Program		Title		Due Date
ASC/CSSE		Cielo capability computing platform production capability readiness		Jun-11
Codes/Simulation Tools Employed: ASC code projects and graphics tools to demonstrate readiness				
Contribution to the ASC Program: The milestone will provide the ASC community with the ability to visualize large-scale simulation datasets coming from the petascale platforms of Roadrunner and Cielo, by acquiring and standing up a visualization compute platform dedicated to this task.				
Contribution to Stockpile Stewardship: The milestone will provide ability to visualize large-scale simulation datasets relevant to stockpile stewardship.				
No.	Risk Description	Risk Assessment (low, medium, high)		
		Consequence	Likelihood	Exposure
1	Loss of key personnel	High	Moderate	Medium
2	Unavailability of selected high-end graphics cards in the milestone time frame	High	Low	Medium
3	Key personnel unavailable because committed to other milestones or high-priority activities	High	Moderate	Medium-High
4	Difficulty in integrating platform because of issues with new technology introduced	High	Moderate	Medium

Milestone (ID#): Demonstration of a legacy application's path to exascale		
Level: 2	Fiscal Year: FY12	DOE Area/Campaign: ASC
Completion Date: March 31, 2012		
ASC nWBS Subprogram: CSSE		
Participating Sites: SNL		
Participating Programs/Campaigns: ASC, DOE Office of Science INCITE program		
Description: Cielo is expected to be the last capability system on which existing ASC codes can run without significant modifications. This assertion will be tested to determine where the breaking point is for an existing highly scalable application. The goal is to stretch the performance boundaries of the application by applying recent CSSE R&D in areas such as resilience, power, I/O and VIZ services, SMARTMAP, lightweight LWKs, virtualization, simulation, and feedback loops. An existing INCITE award on Jaguarpf and/or CCC allocations will be used to quantify the impact of system-level changes to extend the life and performance of the ASC code base. Simulation will be used to supplement the INCITE/CCC calculations at higher scales than are currently available.		

Milestone (ID#): Cielo capability computing platform production readiness		
Level: 2	Fiscal Year: FY12	DOE Area/Campaign: ASC
Completion Date: December 31, 2011		
ASC nWBS Subprogram: CSSE		
Participating Sites: LLNL, LANL, SNL		
Participating Programs/Campaigns: ASC		
Description: The ACES partnership between LANL and SNL is responsible for the deployment and integration of the Cielo Platform that will be sited at LANL. Cielo shall achieve Production Capability Readiness as defined by the Capability Platform Level 2 Milestones Working Group. In summary, this includes the platform is made available for CCC capability work; all system software, tools, utilities and user support processes are available and fully functional; ASC applications targeted for the platform are ported and made available to designers, analysts, and engineers; the platform has demonstrated acceptable reliability performance targets.		

Milestone (ID#): Deploy TLCC2 and a common capacity computing environment	
Fiscal Year: FY12	DOE Area/Campaign: ASC
Completion Date: September 30, 2012	
ASC nWBS Subprogram: CSSE	
Participating Sites: LLNL, SNL, LANL	
Participating Programs/Campaigns: ASC	
Description: Deploy CCE capabilities developed during FY11, including the next major release of common operating system and software stack. Deploy the next generation of the ASC TLCC systems (TLCC2), which will include hardware and software integration and testing for the tri-lab environment.	

VI. ASC Roadmap Drivers for FY11–FY12

One multi-year site objective drives the work in this year's implementation plan:

- Deliver advanced capabilities for QMU and UQ through the Predictive Capability Assessment Project

ASC Roadmap drivers for FY12 also fall within the PCF.

VII. Performance Measures

Table VII-1. ASC Campaign Annual Performance Results (R) and Targets (T)

Performance Indicators	FY07 Results	FY08 Results	FY09 Results	FY10	FY11	FY12	FY13	FY14	FY15	FY16	Endpoint Target
Secretarial Goal: Security: Reduce nuclear dangers and environmental risks GPRA Unit Program Number: 30, Advanced Simulation and Computing Campaign											
Adoption of ASC Modern Codes: The cumulative percentage of simulation runs that utilize modern ASC-developed codes on ASC computing platforms as measured against the total of legacy and ASC codes used for stockpile stewardship activities. (Long-term Outcome) ⁶	R: 63% T : 63%	R: 72% T: 72%	R: 80% T: 80%	T: 85%	T: 90%	T: 95%	T: 100%	N/A	N/A	N/A	By 2013, ASC-developed modern codes are used for all simulations on ASC platforms. Adoption of Modern ASC Codes will enable a responsive simulation capability for the nuclear security enterprise. This measure is meant to show how quickly ASC codes are being adopted by the user community in place of legacy codes.
Reduced Reliance on Calibration: The cumulative percentage reduction in the use of calibration “knobs” to successfully simulate nuclear weapons performance. (Long-term Outcome) ^a	R: 8% T : 8%	R: 16% T: 16%	R: 25% T: 25%	T: 30%	T: 35%	T: 40%	T: 45%	T: 50%	T: 55%	T: 60%	By 2024, 100% of selected calibration knobs affecting weapons performance simulation have been replaced by science-based, predictive phenomenological models. Reduced reliance on calibration will ensure the development of robust ASC simulation tools. These tools are intended to enable the understanding of the complex behaviors and effect of nuclear weapons, now and into the future, without nuclear testing.

⁶ Performance measures were revised in 2007 to be consistent with new program roadmap.

Performance Indicators	FY07 Results	FY08 Results	FY09 Results	FY10	FY11	FY12	FY13	FY14	FY15	FY16	Endpoint Target
ASC Impact on SFI Closure: The cumulative percentage of nuclear weapon Significant Finding Investigations (SFIs) resolved through the use of modern (non-legacy) ASC codes, measured against all codes used for SFI resolution. (Long-term Outcome) ^a	R: 25% T : 25%	R: 37% T: 37%	R: 50% T: 50%	T: 60%	T: 65%	T: 70%	T: 80%	T:85%	T:100%	T:100%	By 2015, ASC codes will be the principal tools for resolution of all SFIs. This demonstrates how valuable the ASC tools are for meeting the needs of the weapon designer's analysts by documenting the impact on closing SFIs.
Code Efficiency: <u>The cumulative percentage of simulation turnaround time reduced while using modern ASC codes. (Efficiency)^a</u>	<u>R: 7%</u> <u>T : 7%</u>	<u>R: 13%</u> <u>T: 13%</u>	<u>R: 13%</u> <u>T: 13%</u>	<u>T: 15%</u>	<u>T: 20%</u>	<u>T: 27%</u>	<u>T: 34%</u>	<u>T: 42%</u>	<u>T: 50%</u>	<u>T: 50%</u>	<u>By 2015, achieve a 50% reduction in turnaround time, as measured by a series of benchmark calculations, for the most heavily used ASC codes. To show code efficiency by demonstrating that simulation time decreases as the ASC codes mature.</u>

VIII. Budget

1.5.1 Integrated Codes						
1.5.1.1	Modern Multi-physics Codes		58.123	0.000	0.000	58.123
		LLNL	30.646			30.646
		LANL	27.477			27.477
		SNL				0.000
		Other				0.000
1.5.1.2	Legacy Codes		0.419	0.000	0.000	0.419
		LLNL	0.419			0.419
		LANL				0.000
		SNL				0.000
		Other				0.000
1.5.1.3	Engineering Codes		31.272	0.000	0.000	31.272
		LLNL	2.833			2.833
		LANL				0.000
		SNL	27.939			27.939
		Other	0.500			0.500
1.5.1.4	Focused Research, Innovation & Collaboration		36.499	0.000	19.500	55.999
		LLNL	14.753			14.753
		LANL	11.609			11.609
		SNL	10.137			10.137
		Other			19.500	19.500
1.5.1.5	Emerging & Specialized Codes		17.745	0.000	0.000	17.745
		LLNL	5.042			5.042
		LANL	7.828			7.828
		SNL	4.875			4.875
		Other				0.000

1.5.2 Physics and Engineering Models						
1.5.2.1	Theoretical Models and Experimental Integration		24.796	0.000	0.000	24.796
		LLNL	9.610			9.610
		LANL	11.331			11.331
		SNL	3.855			3.855
		Other				0.000
1.5.2.2	Model Implementation		27.088	0.000	0.000	27.088
		LLNL	6.953			6.953
		LANL	13.143			13.143
		SNL	6.992			6.992
		Other				0.000
1.5.2.3	Fundamental Physics Codes and Application		6.972	0.000	0.000	6.972
		LLNL	1.706			1.706
		LANL				0.000
		SNL	5.266			5.266
		Other				0.000
1.5.2.4	Material Data Libraries		4.366	0.000	0.000	4.366
		LLNL	1.304			1.304
		LANL	3.062			3.062
		SNL				0.000
		Other				0.000
1.5.2.5	Russian Programs		2.250	0.000	0.000	2.250
		LLNL	0.750			0.750
		LANL	0.750			0.750
		SNL	0.750			0.750
		Other				0.000

1.5.3 Verification and Validation						
1.5.3.1	V&V Methods		23.614	0.000	0.000	23.614
		LLNL	7.800			7.800
		LANL	10.214			10.214
		SNL	5.600			5.600
		Other				0.000
1.5.3.2	Primary V&V Assessments		6.866	0.000	0.000	6.866
		LLNL	4.150			4.150
		LANL	2.716			2.716
		SNL				0.000
		Other				0.000
1.5.3.3	Secondary V&V Assessments		4.011	0.000	0.000	4.011
		LLNL	2.443			2.443
		LANL	1.568			1.568
		SNL				0.000
		Other				0.000
1.5.3.4	Engineering V&V Assessments		15.159	0.000	0.000	15.159
		LLNL	0.522			0.522
		LANL	1.568			1.568
		SNL	13.069			13.069
		Other				0.000
1.5.3.5	Specialized V&V Assessments		2.707	0.000	0.000	2.707
		LLNL	1.200			1.200
		LANL				0.000
		SNL	1.507			1.507
		Other				0.000
1.5.3.6	Data Validation & Archiving		3.035	0.000	0.000	3.035
		LLNL	1.033			1.033
		LANL	1.568			1.568
		SNL	0.434			0.434
		Other				0.000

1.5.4 Computational Systems and Software Environment						
1.5.4.1	Capability Systems	LLNL	2.179	26.600	0.000	28.779
		LLNL	1.092			1.092
		LANL	1.087	26.600		27.687
		SNL				0.000
		Other				0.000
1.5.4.2	Capacity Systems	LLNL	0.000	19.458	0.000	19.458
		LLNL				0.000
		LANL				0.000
		SNL				0.000
		Other		19.458		19.458
1.5.4.3	Advanced Systems	LLNL	4.489	45.942	3.000	53.431
		LLNL	0.100	40.942		41.042
		LANL		2.000		2.000
		SNL	4.389		3.000	7.389
		Other		3.000		3.000
1.5.4.4	System Software and Tools	LLNL	18.572	0.000	0.000	18.572
		LLNL	3.874			3.874
		LANL	7.007			7.007
		SNL	7.691			7.691
		Other				0.000
1.5.4.5	Input/Output, Storage Systems, and Networking	LLNL	17.994	6.000	0.000	23.994
		LLNL	10.560	2.000		12.560
		LANL	2.336	2.000		4.336
		SNL	5.098	2.000		7.098
		Other				0.000
1.5.4.6	Post-processing Environments	LLNL	12.174	3.000	0.000	15.174
		LLNL	2.485	1.000		3.485
		LANL	4.710	1.000		5.710
		SNL	4.979	1.000		5.979
		Other				0.000
1.5.4.7	Common Computing Environment	LLNL	9.333	0.000	0.795	10.128
		LLNL	3.584		0.795	4.379
		LANL	3.102			3.102
		SNL	2.647			2.647
		Other				0.000

1.5.5 Facility Operations and User Support				
1.5.5.1 Facilities, Operations and Communications	LLNL	0.000	0.000	0.000
	LANL			0.000
	SNL			0.000
	Other			0.000
				0.000
1.5.5.2 User Support Services	LLNL	11.363	0.000	12.446
	LANL	3.992		5.075
	SNL	4.229		4.229
	Other	3.142		3.142
				0.000
1.5.5.3 Collaborations	LLNL	5.534	0.000	21.099
	LANL	2.031		2.331
	SNL	0.774		0.963
	Other	1.250		2.255
		1.479	14.071	15.550
1.5.5.4 System and Environment Administration and Operations	LLNL	34.391	46.300	80.691
	LANL	14.591	31.000	45.591
	SNL	12.851	13.000	25.851
	Other	6.949	2.300	9.249
				0.000
1.5.5.5 Facilities, Network and Power	LLNL	6.700	39.404	47.554
	LANL	2.000	10.000	12.000
	SNL	3.000	20.600	23.600
	Other	1.700	5.500	8.400
			3.304	3.554

		PEOPLE	INFOSTRUCTURE	CONTRACTS	Total
	Integrated Codes	144.058	0.000	19.500	163.558
	Physics and Engineering Models	65.472	0.000	0.000	65.472
	Verification and Validation	55.392	0.000	0.000	55.392
	Computational Systems and Software Engineering	64.741	101.000	3.795	169.536
	Facility Operations and User Support	57.988	85.704	18.098	161.790
	Total	387.651	186.704	41.393	615.748
		63%	30%	7%	615.748
		PEOPLE	INFOSTRUCTURE	CONTRACTS	Total
	LLNL	135.473	84.942	2.178	222.593
	LANL	131.930	65.200	0.189	197.319
	SNL	118.269	10.800	5.205	134.274
	Other	1.979	25.762	33.821	61.562
		387.651	186.704	41.393	615.748
			Month	Monthly Cost (\$Ms)	Cumulative Cost (\$Ms)
			Oct	28.510	28.510
			Nov	46.119	74.629
			Dec	44.090	118.719
			Jan	28.485	147.204
			Feb	34.791	181.995
			Mar	44.321	226.316
			Apr	37.212	263.528
			May	38.036	301.564
			Jun	49.810	351.374
			Jul	57.473	408.847
			Aug	54.136	462.983
			Sep	112.006	574.989

Appendix A. Glossary

1D	One Dimensional
2D	Two Dimensional
3D	Three Dimensional
AAAS	Annual Archive Allowance System
ACES	New Mexico Alliance for Computing at Extreme Scale
ADEPT	Applications Development Environment and Performance Team
AMR	Adaptive Mesh Refinement
API	Application Programming Interface
ASC	Advanced Simulation and Computing
ASCI	Accelerated Strategic Computing Initiative
CAD	Computer Aided Design
CCC	Capability Computing Campaign
CCE	Common Computing Environment
CD	Critical Decision
CRASH	Center for Radiative Shock Hydrodynamics
CSSE	Computational Systems and Software Environment (WBS 1.5.4)
D&E	Development and Engineering
DDS	Distributed Data Services
DoD	Department of Defense
DOE	Department of Energy
DSL	Domain-Specific Language
DTRA	Defense Threat Reduction Agency
EOS	Equation of State
FPGA	Field Programmable Gate Arrays
FOUS	Facility Operations and User Support (WBS 1.5.5)
FVM	Finite Volume Method
GA	General Availability
GB	Gigabytes
GB/sec.	Gigabytes per Second
GPFS	Global Parallel File System

GPGPU	General-Purpose Computing on Graphics Processing Units
HPA	Host Processing Agent
HTGL	High Temperature Gasdynamics Laboratory
GPA	Guest Processing Agent
HPC	High Performance Computing
HPSS	High-Performance Storage System
HQ	ASC Headquarters
I/O	Input/Output
IAA	Institute for Advanced Architectures and Algorithms
ID	Initial Delivery
IDA	Interactive Data Analysis
IHPC	Inter-Site High Performance Computing
ISA	Tri-Labs Interconnection Security Agreement
JEI	Joint Exascale Initiative
LAN	Local Area Network
LANL	Los Alamos National Laboratory
LC	Livermore Computing Center
LDCC	Laboratory Data Communications Center
LEP	Life Extension Program
LLNL	Lawrence Livermore National Laboratory
LVOC	Livermore Valley Open Campus
LWK	Lightweight Kernel
MPI	Message Passing Interface
NAPS	NNSA Policy Letters (known as...)
NAS	Network-Attached Storage
NFS	Network File System
NIF	National Ignition Facility
NNSA	National Nuclear Security Administration
NPR	Nuclear Posture Review
NSA	National Security Agency
nWBS	National Work Breakdown Structure
O SS	Open SpeedShop
OCF	Open Computing Facility
OGA	Other Government Agency
PCF	Predictive Capability Framework

PECOS	Predictive Engineering and Computational Sciences
PLFS	Parallel Log Structured File System
PRISM	Prediction of Reliability, Integrity, and Survivability of Microsystems (Center for...)
PSAAP	Predictive Science Academic Alliance Program
PSI	Parallel Storage Interface
PUE	Power Usage Effectiveness
QMU	Quantification of Margins and Uncertainties
QPX	Quad Floating Point
R&D	Research and Development
RAID	Redundant Array of Independent Disks
RAIT	Redundant Array of Independent Tapes
RANS	Reynolds-Averaged Navier Stokes
RFP	Request for Proposal
RHCE	Resource Health Characterization Engines
RHEL	Red Hat Enterprise Linux
SAN	Storage Area Network
SAP	Scalable Applications Preparations
SARAPE	Synchronized Account Request Automated ProcEss
SCC	Nicholas C. Metropolis Center for Modeling and Simulation
SCF	Secure Computing Facility
SCR	Scalable/Check-Point Restart Code
SFI	Significant Finding Investigation
SLA	Service Level Agreement
SLURM	Simple Linux Utility for Resource Management
SNL	Sandia National Laboratories
SSP	Stockpile Stewardship Program
SST	Structural Simulation Toolkit
SU	Scalable Unit(s)
TBI	Thermonuclear Burn Initiative
TLCC	Tri-Lab Linux Capacity Cluster
TLS	Thread Level Speculation
TM	Transactional Memory
TOSS	Tripod Operating System Software
TSF	Terascale Simulation Facility
UQ	Uncertainty Quantification

V&V	Verification and Validation
VTK	Visualization Toolkit
WAN	Wide Area Network
WC Tool	Workload Characterization Tool

Appendix C. Points of Contact

WBS	Title	Contact
1.5.4	Computational Systems and Software Environment	Becky Springmeyer, LLNL, 925-423-0794, springmeyer1@llnl.gov John Thorp, LANL, 505-665-82265, thorp@lanl.gov Sudip Dosanjh, SNL, 505-845-7018, ssdosan@sandia.gov
1.5.4.1	Capability Systems	Mark Seager, LLNL, 925-423-3141, seager1@llnl.gov Manuel Vigil, LANL, 505-665-1960, mbv@lanl.gov Sudip Dosanjh, SNL, 505-845-7018, ssdosan@sandia.gov
1.5.4.2	Capacity Systems	Mark Seager, LLNL, 925-423-3141, seager1@llnl.gov Manuel Vigil, LANL, 505-665-1960, mbv@lanl.gov Sudip Dosanjh, SNL, 505-845-7018, ssdosan@sandia.gov
1.5.4.3	Advanced Systems	Mark Seager, LLNL, 925-423-3141, seager1@llnl.gov John Thorp, LANL, 505-665-8226, thorp@lanl.gov Sudip Dosanjh, SNL, 505-845-7018, ssdosan@sandia.gov
1.5.4.4	System Software and Tools	Becky Springmeyer, LLNL, 925-423-0794, springmeyer1@llnl.gov John Thorp, LANL, 505-665-8226, thorp@lanl.gov Sue Kelly, SNL, 505-845-9770, smkelly@sandia.gov
1.5.4.5	I/O, Storage Systems, and Networking	Becky Springmeyer, LLNL, 925-423-0794, springmeyer1@llnl.gov John Thorp, LANL, 505-665-8226, thorp@lanl.gov Ron Oldfield, SNL, 505-284-9153, raoldfi@sandia.gov
1.5.4.6	Post-Processing Environments	Becky Springmeyer, LLNL, 925-423-0794, springmeyer1@llnl.gov John Thorp, LANL, 505-665-8226, thorp@lanl.gov David Rogers, SNL, 505-844-5323, dhroger@sandia.gov
1.5.4.7	Common Computing Environment	Becky Springmeyer, LLNL, 925-423-0794, springmeyer1@llnl.gov John Thorp, LANL, 505-665-8226, thorp@lanl.gov Judy Sturtevant, SNL, 505-845-9448, jesturt@sandia.gov
1.5.5	Facility Operations and User Support	Kim Cupps, LLNL, 925-423-7262, cupps2@llnl.gov Hal Armstrong, LANL, 505-667-8426, hga@lanl.gov John Zepper, SNL, 505-845-8421, jdzepper@sandia.gov

WBS	Title	Contact
1.5.5.1	Facilities, Operations, and Communications	Kim Cupps, LLNL, 925-423-7262, cupps2@llnl.gov Hal Armstrong, LANL, 505-667-8426, hga@lanl.gov John Noe, SNL, 505-844-5592, jpnoe@sandia.gov Jim Mahoney, KCP, 816-997-5158, jmahoney@kcp.com
1.5.5.2	User Support Services	Kim Cupps, LLNL, 925-423-7262, cupps2@llnl.gov Hal Armstrong, LANL, 505-667-8426, hga@lanl.gov Dino Pavlakos, SNL, 505-844-9089, cjpavla@sandia.gov
1.5.5.3	Collaborations	Brian Carnes, LLNL, 925-423-9181, carnes1@llnl.gov Hal Armstrong, LANL, 505-667-8426, hga@lanl.gov Kenneth Alvin, SNL, 505-844-9329, kfalvin@sandia.gov
1.5.1.4-TRI-001	Caltech, Center for Simulating Dynamic Response of Materials	Michael Ortiz, 626-395-4530, Ortiz@aeor.caltech.edu Mark Stalzer, 626-395-2521, stalzer@caltech.edu Susan Powell, 626-395-2909, spowell@cacr.caltech.edu
1.5.1.4-TRI-002	Purdue University, Center for Prediction of Reliability, Integrity and Survivability of Microsystems	Jayathi Murthy, 765-494-5701, jmurthy@ecn.purdue.edu Dawn Weisman, 765-494-8336, dweisman@purdue.edu
1.5.1.4-TRI-003	Stanford, Center for Integrated Turbulence Simulation	Parviz Moin, 650-723-9713, moin@stanford.edu Gianluca Iaccarino, 650-723-9599, iops@stanford.edu Deb Michael, 650-725-2077, debmic@stanford.edu
1.5.1.4-TRI-004	University of Michigan, Center for Radiative Shock Hydrodynamics	Paul Drake, 734-763-4072, rpdrake@umich.edu James Holloway, 734-936-3126, hagar@umich.edu
1.5.1.4-TRI-005	University of Texas, at Austin, Center for Predictive Engineering and Computational Sciences	Bob Moser, 512-471-0093, rmoser@mail.utexas.edu Chris Simmons, 512-232-2881, csim@colab.ices.utexas.edu

Appendix D.

WBS 1.5.1.4-TRI-001 Academic Alliance Centers

The Academic Alliance Centers project includes research activities at the funded academic centers as part of the PSAAP program, as listed below.

- California Institute of Technology (Caltech)
- Purdue University
- Stanford University
- University of Michigan
- University of Texas, at Austin

California Institute of Technology

The Center for the Predictive Modeling and Simulation of High-Energy Density Dynamic Response of Materials

Caltech's Multidiscipline Simulation Center's overarching objective is the development of a multidisciplinary predictive science methodology focusing on high-energy-density dynamic response of materials and the demonstration of the methodology by means of a concerted and highly integrated experimental, computational, and analytical effort focusing on an overarching ASC-class problem: hypervelocity normal and oblique impact of projectiles on metallic and non-metallic targets, at velocities up to 10 km/sec. Hypervelocity impact gives rise to pressures in the Mbar range and strain-rates up to 10^{11} /sec., providing a grand-challenge problem in predictive science that is also well-matched to the direct interests of the NNSA mission. The overarching hypervelocity impact application, in conjunction with a rigorous and novel methodology for model-based UQ, will provide the intellectual backbone of the Center and its chief organizing principle. In particular, the QMU will drive and closely coordinate the experimental, computational, modeling, software development, and V&V efforts within a yearly assessment format.

Planned activities in FY11:

- Perform Lagrangian Optimal Transportation Meshfree (OTM) and Eulerian simulations of oblique hypervelocity impact of single steel/aluminum plates and nylon projectiles in the 5-10 Km/sec. impact velocity range, including verification, performance analysis, and UQ runs using the UQ pipeline. (This is an overarching, year-three activity.)
- Perform hypervelocity tests on single steel/aluminum target plates and nylon projectiles in the 5-10 Km/sec. range, including experimental setup and UQ tests. (This is an overarching, year-three activity.)

A summary of year-three activities follow:

- Perform nylon/steel/aluminum hypervelocity Eulerian and OTM simulations for UQ analysis
- Perform nylon/steel/aluminum hypervelocity tests for UQ analysis

- Develop further the UQ pipeline as a heterogeneous and distributed simulation and optimization framework
- Implement into UQ pipeline of optimal-UQ probability-of-failure bounds and of the legacy UQ protocol
- Integrate new strength, EOS, and transport models into the simulation codes
- Implement multispecies plasma simulation capability
- Analyze high-pressure EOS data for steel and aluminum from converging-shock tests
- Develop an eFF methodology for higher Z elements

Preliminary planned activities in FY12:

- Run Eulerian hydrocode simulation of oblique hypervelocity impact of single Ta plate by Ta projectile in the 5–10 km/s impact velocity range
- Implement full-system Ta/Ta hypervelocity runs for UQ analysis, including verification (nonlinear sensitivity analysis) and validation (in coordination with full system experiments) runs
- Deploy improved concentration-of-measure uncertainty bounds that exploit the hierarchical and multiscale structure of hypervelocity
- Deploy a robust UQ-pipeline on a variety of NNSA platforms that supports the VTF and Eulerian simulators
- Implement and validate a fast multiscale model of polycrystalline behavior for Fe, including solid-solid phase transitions
- Study thermal and electrical conductivities and optical frequency response function of the dense plasma state of Ta for a wide range of temperatures and densities

Purdue

Center for Prediction of Reliability, Integrity, and Survivability of Microsystems

The overall objective of PRISM is to accelerate substantially the integration of MEMS technologies into civilian and defense systems. PRISM aims to significantly improve understanding of the long-term reliability of MEMS and survivability in harsh environments by simulating rigorously, and at multiple scales, the physics of failure, accounting for the coupled electrical, mechanical, thermal, and materials behavior of MEMS, from atoms to devices. Advanced simulation software developed by the Center will be encapsulated in an integrated simulation system, MEMOSA.

Planned activities in FY11

- Continue the development of MEMOSA to address 1) structural response under the action of fluidic and electrostatic forces with metal-dielectric contact; 2) continued refinement of dielectric charging, including contact effects; 3) continued development and refinement of mesoscale metal-dielectric contact models informed by molecular dynamics simulations of metal-dielectric contact; 4) continued development and refinement of materials models; and 5) development and refinement of ES-BGK model for rarefied gas dynamics

- Simulate integrated fluid-structure-electrostatic response of the PRISM device, including UQ
- Simulate dielectric charging of metal-insulator-metal (MIM) capacitor for single and multiple voltage cycles, including UQ
- Complete implementation, V&V of ES-BGK model in MEMOSA-finite volume method (FVM)
- Complete testing and benchmarking of integrated parallelized MEMOSA FVM and material point method solvers
- Complete fluid-structure-electrostatics UQ experiments to collect probability density function data for validation; complete dielectric charging and electro-thermal experiments to provide validation data; complete creep experiments to provide calibration data for creep model
- Perform experimental micro-metrology to obtain improved characterizations of the geometry and material properties of the PRISM device

Expected deliverables in FY11:

- Integrated system simulations of the fluid-structure-electrostatic response of the PRISM device, including UQ, for single voltage cycle under normal operating conditions
- Simulation of dielectric charging of metal-insulator-metal (MIM) capacitor for single and multiple voltage cycles, including UQ
- UQ of PRISM device multi-cycle simulations using coarse-grained model
- Mesoscale model for metal-dielectric contact informed by atomistic contact simulations
- Material model for creep in thin metallic membranes
- Data from fluid-structure-electrostatics UQ experiments providing probability density function data for validation; calibration experiments to determine constants in dielectric charging model using MIM capacitor; validation experiments of dielectric charging in PRISM device; and creep experiments to provide calibration data for creep model
- Parallel scaling benchmarks for integrated MEMOSA-FVM and MPM solvers

Preliminary planned activities in FY12:

- Integrate system simulations of the fluid-structure-electrostatic response of the PRISM device, including UQ for multiple voltage cycles under normal operating conditions
- Integrate system simulations of the integrated creep and dielectric charging response of the PRISM device under an applied voltage including UQ under normal operating conditions

Stanford University

The Center for Predictive Simulations of Multi-Physics Flow Phenomena with Application to Integrated Hypersonic Systems

The objective of the Center is to characterize the operability limits of an air-breathing hypersonic vehicle and associated propulsion system using predictive multi-physics

simulations. The primary focus will be on the “unstart” failure mode triggered by thermal choking. Air-breathing hypersonic vehicles are envisioned as a means for reliable low-cost access to space. These vehicles are highly integrated systems whose performance depends on complex physics and the interactions between all of its components. Current state-of-the-art simulation capabilities cannot predict these systems reliably, particularly near their operability limits.

World-class experimental facilities in Stanford's High Temperature Gasdynamics Laboratory (HTGL) will be used to conduct tightly integrated validation experiments for the key component physics and models. The development and implementation of UQ methods for very large systems will be an integral part of the effort at both the component and the system levels. Novel verification methods for high-fidelity simulations will also be developed and implemented. Stanford will leverage advanced computer science methods developed at Stanford to directly impact simulation tools and ensure scalability, program correctness, and portability to future platforms with very large numbers of cores.

Planned activities in FY11:

- Estimate unstart probability for flight conditions using 2D/3D RANS models validated with DLR experiment, jet in supersonic cross-flow experiment, shock/turbulent boundary layer baseline experiment at Stanford
- Extend adjoint solver for mixing and combustion, further develop goal oriented mesh adaptation and integrate adjoint verification framework within simulation suite
- Complete LES of non-reacting and reacting jet in cross-flow cases in Charles to investigate mixing and combustion, particularly inside wall boundary layer, where Joe modeling approaches have deficiencies; perform non-reacting high-fidelity and low-fidelity simulations of shock-boundary layer interactions using the HTGL reflected-shock/boundary layer experimental data sets
- Assess and improve URANS predictions of shock/turbulent boundary layer interaction using LES and DNS data; characterize epistemic uncertainties linked to turbulence model
- Investigate and characterize ignition and flame structure in newly developed scramjet model combustor; use new model combustor, jet in supersonic cross-flow experiment and DLR data (two meetings with DLR planned this year and next year) to validate and refine combustion model for LES and RANS
- Assess radiation importance in combustor and potentially integrate radiation and heat transfer to the walls into RANS solver
- Further characterize unstart process due to mass injection for different flow, jet and boundary layer conditions and develop PIV measurement infrastructure; use knowledge acquired to validate solvers and to better characterize unstart bounds
- Implement fast direct solvers using the Liszt interface, thus, taking advantage of the knowledge of the geometric topology associated with a matrix (for example, for finite-element or volume calculation) to reduce the cost of direct solvers. Port and optimize Sequoia on lab systems including Cerrillos

Expected deliverables in FY11:

- Full system QMU analysis with multiple gates and epistemic uncertainties linked to the turbulence models

- Demonstration of solution verification (error estimation, goal oriented mesh adaption) within a UQ framework in shock/boundary layer interaction and mixing in jet in cross-flow
- Compressible LES solution of non-reacting and reacting jet in cross-flow with experimental comparison
- Reaction rate measurements of selected $\text{HO}_2/\text{H}_2\text{O}_2$ reactions, benchmark studies of unstart in model inlet and validation data for mixing in jet in supersonic cross-flow and combustion in model combustor
- Fast direct solver implemented in Liszt infrastructure

Preliminary planned activities in FY12

- Apply verification framework to overarching problem of system failure characterization within global UQ strategy
- Perform validation with UQ of RANS and LES based on shock/turbulent boundary layer with perturbations and jet in supersonic cross-flow
- LES of unstart in isolator with near-wall models and estimation of epistemic uncertainties for shock-induced separation
- Continue coupled high-fidelity (LES, Charles) and lower-fidelity (RANS, Joe) simulation activity at the sub-system level to characterize the model-form uncertainties in RANS for hypersonic systems, including mixing and combustion

University of Michigan

The Center for Radiative Shock Hydrodynamics (CRASH) is advancing predictive science in the nationally important area of radiation hydrodynamics via a unified, multi-prong approach. To substantially improve the ability to perform predictive simulations of high-energy-density and astrophysical flows, Center researchers:

- Continue to develop software for radiation hydrodynamics to serve as a test bed for development and V&V of radiation hydrodynamics modeling elements
- Continue to develop a system for extensive V&V of the software
- Extended an existing experimental effort, centered on radiative shocks, to quantify uncertainties in the experiments and obtain data focused on improving the predictive capability of the code
- Simulate this sequence of experiments and assess the predictive capability of the simulations
- Established a doctoral program track for predictive science and engineering

Planned activities in FY11:

- Simulate variations on the year five-year experiment in 3D with version 2 of the CRASH code
- Assess predictive capability for year three and year five experiments; this will involve screening for variable selection, other forms of dimension reduction, simulations over distributions of values, including variations reflecting uncertainty in constants of nature, a Bayesian analysis involving a comparison of simulation results to distributions of actual data, and other steps

- Perform experiments to quantify the time evolution of the Xe radiative shock, and make other measurements determined to have utility for assessing that predictive capability of the code
- Modify the BATSRUS/CRASH code to: 1) employ a multilevel preconditioner, 2) have improved internal parallel I/O, 3) enable improved visualization of the output, and 4) add flux limiting to the electron heat conduction
- Further enhance PDT's capability: 1) implement advanced transport-sweep logic and expand scaling studies to tens of thousands of cores; and 2) implement advanced iterative algorithms to achieve significant speedups on CRASH-like problems
- Quantitatively compare solutions of various radiation treatments and spatial resolutions on CRASH-like problems, paying particular attention to energy deposition in plastic walls; the radiation treatments will include: 1) gray diffusion (CRASH); 2) multigroup diffusion (CRASH) with various group structures; 3) multigroup transport (PDT) with various group structures; use the results to set the direction for the combined application of CRASH and PDT
- Continue to work actively on the underlying physics of the radiative shocks of interest, to better develop the expert judgment that is essential for assessing predictive capability
- Continue to develop, implement, and document test and verification problems ranging from unit tests to multi-physics tests to code comparisons to validation against experiments

Expected deliverables in FY11:

- A report of a 3D simulations of the signature experiment with version 2 of the CRASH code, as part of the annual technical report
- A software report, as part of the annual technical report, detailing the accomplishments in code development
- A UQ report, as part of the annual technical report, describing the results of assessment of predictive capability studies
- An experiments report, as part of the annual technical report, describing the analysis of the FY09 and FY10 experiments and the design of the FY11 experiment
- A discussion of progress in fundamental understanding and semi-analytic modeling of the physical system, as part of the annual technical report

Preliminary planned activities in FY12:

- Improve the capabilities of the simulation, with priorities set by what is learned in the assessment of predictive capability
- Identify, plan, and perform experiments based on areas of maximum importance determined by the assessment of predictive capability
- Continue to improve the fundamental understanding of the physical system of interest
- Assess predictive capability for all FY12 experiments and year five experiments

University of Texas

The Center for Predictive Engineering and Computational Sciences

The goal of the PECOS Center is to develop next generation advanced computational and UQ methods for predictive simulation of multi-scale, multi-physics phenomena relevant to the NNSA, and to apply these methods to the problem of reentry of vehicles into the atmosphere.

Simulation of vehicle reentry into the atmosphere requires modeling of the interaction of extremely high temperature gas flows with the high temperature response of materials, in particular the vehicle's thermal protection system. The high gas temperatures produce chemical dissociation, thermal non-equilibrium, and possibly ionization. Radiative heat transfer is an important part of the heat load on the vehicle, while transition and turbulence greatly enhance the rate of heat transfer. During reentry, the ablative thermal protection system responds via pyrolysis, chemical reaction, and formation and mechanical degradation of a refractory char layer. Models of these high-energy, multiscale, multiphysics phenomena will be integrated into a unified simulation code, designed to support predictive simulation.

Planned activities in FY11:

- Integrate FIN-S with updated physics models in radiation and turbulence
- Implement verifiable high temperature non-equilibrium code
- Perform initial simulation of high-speed boundary layers via DNS
- Develop efficient variants of Stochastic Newton method
- Explore information theoretic driven, QoI-aware optimal experimental design

Expected deliverables in FY11:

- Completion of first propagation of uncertainty in the full system simulation using best available single physics calibrations
- Completion of adjoint-enabled hypersonic flow code with coupled multiphysics
- Completion of single physics calibrations across all modeling domains

Preliminary planned activities in FY12:

- Complete first propagation of uncertainty in the full system simulation using validated single physics and calibrated multiphysics
- Continue working on model development, implementation, and code verification
- Continue working on model validation using updated data across all domains
- Develop adjoint-enabled modeling codes across all domains

Appendix E. ASC Obligation/Cost Plan

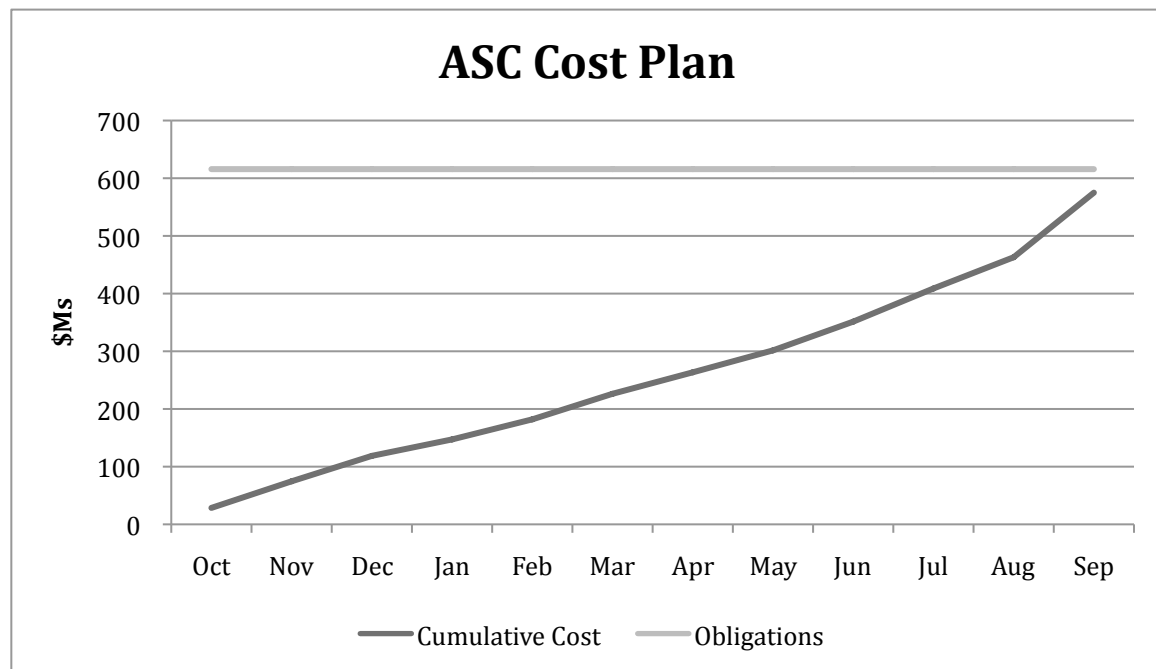


Figure E-1. ASC obligation/cost plan for FY11.

